LXF121
LXF August 2009

OpenOffice.org » GP2X » Code

LINUX FORMAT

The #1 source for Linux

LXFDVD

Linux Mint 7

Full distro: Ubuntu taken to the freshest next level

PLUS: OpenSolaris 2009.06 » Ulteo 1.0 » OOo 3.1

Pick the perfect netbook

8 Low-cost, Linux-ready laptops tested – find the best one for you!

Linux meets Mac

How to put your Linux skills to work on Mac OS X p50

OpenOffice.org 3.1

Come in Word, your time is up p32

26 pages of great tutorials inside

Stream music

Play tunes over your home network p90

Talk to the world

Flex your coding fingers with an awesome chatbot p94

Secure servers

Nessus + iptables = safe boxes on the net p98

Netbooks Mega-Roundup » Mac OS X » Benchmarking » Sox » Jets'n'Guns » Python Hacking » Michael Simms » Ulteo 1.0 » ↑↓←→ ABC + Start

"Every game we've had was delayed in one way or another – I won't let a bad game out"

Linux Game Publishing's Michael Simms p48

techradar.com

LXF121 AUGUST 2009 £6.49

Future MEDIA WITH PASSION

9 771473 968012 08

www.linuxformat.com

Printed in the UK

UKFAST®

// speed through innovation

cheetah clients enjoy

300%

growth and more...

Some hard and fast results from our new dedicated hosting clients:

- // **300%** increase in sales in first 7 days (www.topfurniture.co.uk)
- // **300%** increase in page load speed (www.foresitespa.com)
- // **300%** increase in site traffic within 4 weeks (www.alljoinon.com)
- // Google page 1 ranking within 6 weeks of launch (www.letsstaytogether.co.uk)

If fast results are important to your business, then don't hesitate, call us today.

0800 954 0894 **call now**

www.ukfast.co.uk/results.html

The ISPAs
05/06/07/08 WINNER
Best Hosting Provider

LINUX FORMAT
The #1 source for Linux

What we do

» We support the open source community by providing a resource of information, and a forum for debate.

» We help all readers get more from Linux with our tutorial section – we've something for everyone!

» We license all the source code we print in our tutorials section under the GNU GPLv3.

» We give you the most accurate, unbiased and up-to-date information on all things Linux.

Who we are

After the netbook fest on 38, there's only one thing we want to know: if you could add just one killer feature to your ultimate netbook, what would it be?

Graham Morrison
"I would replace the webcam with a machine that could detect involuntary dilation of the iris."

Mike Saunders
"Beer-based ethanol fuel-cells. One for you, one for me... (hic)."

Neil Bothwick
"I'd protect its limited resources by only allowing three programs to run at a time. Hold on a mo, can I patent that idea?"

Efrain Hernandez-Mendoza
"Phased plasma pulse rifle in 40-watt range. 45 longslide with laser sighting. Uzi 9mm."

Andrew Gregory
"An applet that measures how much booze you drink and calculates how much Marmite you need to eat to ward off Korsakoff's."

Andy Channelle
"I wouldn't add anything. I'd just change the name to something more catchy, like: Low-Cost Small Notebook."

David Cartwright
"Fingerprint recognition technology. Failing that, a laser beam with a shark on its head."

Andy Hudson
"Some sort of geolocation device so I can find the darn thing when my missus 'tidies' it away."

Nick Veitch
"An emergency inflatable desk/chair/secretary combo for those long waits in airports."

Susan Linton
"They should include a swing-out bottle opener and cork screw – like a Swiss Army netbook."

Shashank Sharma
"Netbooks should have an arcade-style coin slot for guest users. The collected money goes into the buy-useful-machine fund."

Mayank Sharma
"Netbooks should be able to bend the laws of physics to slow time, especially when working on processor-intensive tasks."

Netbooks galore!

» On a normal PC, the cost of OEMs throwing in a copy of Windows Vista as standard is pretty small compared with the overall cost of the machine. But in the world of netbooks, where prices almost sneak under the £100 mark, even a super-discounted copy of Windows will eat up 20% or more of the system's cost, which gives Linux a decided edge. Yes, Linux's free-as-in-speech nature is as important as ever, but in this particular occasion being free as in beer is what matters to most people. If they choose Linux they save money – often enough to upgrade to a more expensive netbook.

But which one to choose? It's very easy to think that all netbooks are made alike, but there's more to it than that: there are several different Linux distros for netbooks, and some don't even come with Linux at all. Either way, if you choose to remove the built-in distro and try something else, how good are other distros at recognising the hardware, and just what are the distro alternatives for netbooks? All these things and more are covered in our feature starting on p38.

STOP PRESS

As part of our continued efforts to reduce environmental waste, this is the first issue of Linux Format to come with an Eco Disc rather than a standard DVD. We think these are cool because you can bend them 180 degrees, but we hope you'll like their extra durability – and whale-saving abilities!

Paul Hudson Editor
» paul.hudson@futurenet.co.uk

LINUX FORMAT

Contents

Everything that's inside this packed issue of the world's finest Linux publication.

Reviews

GP2X Wiz 24
The second attempt at an open-ended handheld console is here.

Kdenlive 26
Video editing has long been a prickly point in free software, but this finally fixes it.

Ulteo OVD 28
You'll no doubt remember or talk of an online distro and thought we were crazy, but Gaël Duval thinks different...

› **The Wiz is Gamepark's second major attempt at handheld gaming.**

Jets'n'Guns 29
Shoot, shoot, then shoot some more. If there is a plot to this glorious – ahem – tribute to *R-Type*, we don't need it!

EnergyXT 2.5 30
Before MP3s there were MODs, and now you can add beat slicing and more!

OpenOffice.org 3.1 32
It's prettier and has more features, but don't get your hopes up: it's still slower than a snail dipped in treacle.

Books 34

› **Books: good for the mind. And for keeping doors open.**

Netbook megatest!

We put Linux onto eight top netbooks then test them until they beg for mercy p38

What on Earth is...?

APML

Tell websites what you care about and skip the dross p60

Benchmarking p56

Make your code run faster with just a little tweakage.

Talking heads

"There needs to be a clear demarcation of quality when it comes to major version releases"

Our Graham on versioning vagaries p23

Linux Mint 7

Fresh new Ubuntu respin

» **Ulteo** User-friendly gets taken further
» **OpenSolaris** Sun's finest is here to try
» **OpenOffice.org 3.1** Upgrade today
... plus our TuxRadar podcasts and more! p66

Don't miss...

Sox audio editor 54
Make changes to audio files in seconds, and straight from the command line too!

Mac OS X..... 50
Put your hard-earned Linux skills to work on strangely shiny hardware.

Kill spam for good..... 90
Unless you like collecting email adverts for dodgy Viagra pills, use free software to kill spam dead!

Regulars at a glance

Newsdesk 6
Microsoft and Linux in force-joining shocker, Red Hat takes on Oracle and Novell's profits are up.

User groups..... 16
Written by user groups for user groups, these are all yours!

Mailserver..... 18
It's not just the police, members of parliament and judges who are getting younger – *Linux Format* readers are too.

Distrowatch..... 36
Yes, Mandriva and Ubuntu snatch the headlines, but there's room to discuss SliTaz and your monthly dose of distro statistics.

What on Earth..... 60
Letting websites know what you like sounds like a privacy nightmare, but read a little more about APML and we think you'll like the sound of it!

Sysadmin..... 62
Have all your redshirts been shot on away missions? Let SSH and *Wireshark* take their place!

HotPicks 70
These are so hot we had to make the rest of the magazine out of liquid nitrogen to balance things out.

Subscriptions .. 76/102
Get issues sent straight to your door, save money and help keep the LXF staff off the streets.

Answers 105
YOUR PROBLEMS SOLVED
This month, we went to the pub first and then answered your questions.

Next month 113
LXF121 will be in your grubby mitts faster than a speeding hedgehog.

› **Our subscriptions team is waiting for your call on p48.**

Tutorials

Beginners
Basket and Sockso..... 78
Don't let the idea of building a music server fill you with dread – we make it easy!

› **Share your music with the world, or your bathroom – you're in control!**

Gimp
She's on fire!..... 82
Feeling artistic? Set your phasers to win and try out this great *Gimp* tutorial.

› **Setting your hair on fire: a rather extreme way of beating dandruff.**

OpenOffice.org
Script quizzes 86
Fact: multiple-choice tests can be generated on the fly with *OOo*. Learn how here!

Networking
Get rid of spam..... 90
Setting up bulletproof spam filters is just one more reason why *Postfix* rocks.

Python
Code a chatbot..... 94
Learn how Python can build a chatbot that says, "what?" and "where's the tea?".

Hardcore Linux
Server security 98
Nessus, *Tripwire* and more – with their powers combined you are Captain Security!

COMMUNITY

MS joins Linux Foundation

Unlikely bedfellows get together to battle warranty changes.

Microsoft and the Linux Foundation, two organisations usually guaranteed to be on different sides of any argument, have joined forces. The unlikely union was formed to combat legal changes in the US that will mean software has an implied fit for purpose warranty. The two organisations drafted a letter together to the American Law Institute (ALI) that expressed their concerns about a new set of principles on the laws surrounding software contracts.

The ALI is an organisation of legal experts that constructs restatements of US law with the intention of convincing judges to adopt their interpretations when deciding cases. The concern of the Linux Foundation and Microsoft is that the information that feeds into this process is assembled behind closed doors and without external input.

In the letter, lawyers for MS and the Linux Foundation said the principles had been written without consultation with software developers, and that the changes would lead to disruption of a well-functioning market and increase the risk of litigation. The letter said:

› The header is odd, but the words contain good sense.

› Microsoft's Horacio Gutierrez famously said Linux violated hundreds of MS patents. Now he's writing a letter with the Linux Foundation.

"Notwithstanding our approaches to the licensing and distribution of software, we share a common desire for a sound, effective commercial law framework for software that reflects business and community realities."

Linux Foundation chief executive Jim Zemlin also said that the new principles would: "interfere with the natural operation of open source licences and commercial licences".

> **"There's a wide range of issues that affect all software developers."**

The situation regarding free software, which may be given away or sold, is particularly problematic. The letter said: "One can download Ubuntu free of charge over the internet. But one can also buy Ubuntu on CD or DVD. Thus the warranty may or may not apply to the exact same software distribution depending upon how the recipient obtains it."

The proposals causing the biggest headache for both the Linux Foundation and Microsoft are two non-disclaimable warranties that are implied whenever an individual or organisation distributes software. The first requires the program not to infringe any patents and copyrights, which isn't an easy task since a developer would have to have in-depth knowledge of all valid patents.

The second requires software to be free from material defects. This could lead to open source developers being sued for producing software with bugs if it's ever incorporated in a commercial product. The subject is covered in a section of the principles that say: "A transferrer who receives money or a right to payment ... in exchange for the software warrants to any party in the normal chain of distribution that the software contains no material hidden defects of which the transferrer was aware at the time of the transfer."

Licence clashes

The Linux Foundation said this conflicts with many open source licences, not least the GPL, which states software is provided "as is without warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose". Moreover, Microsoft argues that ALI fails to take into account the complexity of software distribution, where the traditional proprietary and free models have been joined by a class of software that's delivered as a service.

Horacio Gutierrez, Microsoft VP and deputy general counsel, wrote in his blog: "The mere fact that the Linux Foundation and Microsoft are joining forces may be viewed by some as remarkable, given that our differences receive far more public attention than when our interests converge. But there is a wide range of issues that affect all software developers alike."

ENTERPRISE

Red Hat and Ingres go on the offensive

Oracle's purchase of Sun Microsystems galvanises the competition to start working together.

Since the purchase of Sun Microsystems by Oracle, Red Hat and Ingres have developed an increasingly cosy relationship that has now resulted in the launch of an Ingres Development Stack. The package bundles together an Ingres database and *JBoss Developer Studio* running on the *JBoss Application Server*, which in turn runs on top of Red Hat Enterprise Linux 5.

Roger Burkhardt, CEO of Ingres, said that as IT budgets begin to fall, companies of all sizes are starting to look at the licence fees they pay to developers of operating systems, middleware and databases. "Globally, IT organisations spend upwards of $10bn each year on licence fees," he said, "and have found that they can cut costs and still innovate if they deploy on a full open source infrastructure stack."

Reassessed costs

The savings for larger enterprises appear to be significant. A standard annual support contract for the Ingres database deployed on commodity hardware is likely to cost in the region of $39,000, although this doesn't include Red Hat's support costs. A similar setup based on Oracle may cost up to $700,000. But while Ingres is keen to push the Red Hat angle, the software is also certified to run on Ubuntu and SUSE Linux Enterprise Server.

Financial software developer Biveroni Batschelet Partners AG (BBP) is one of the early adopters of the system. The company's CTO, Amir Housseini, said BBP would use the new stack to speed up inter-bank connections. "Financial institutions place their absolute trust and the success of their day-to-day business processes directly in our hands," he said. "Given the sensitivity of our application and the underlying data, we require an application infrastructure that can meet extremely stringent requirements, especially for security and availability."

Ingres CEO Roger Burkhart stressed the benefits of an open source infrastructure.

Oracle/Sun watch

Oracle's acquisition of Sun appears to be running smoothly, but speculation is growing that the merged company may attempt to lose Sun's hardware business to concentrate fully on software, especially as it puts Oracle on a potential collision course with partners such as Dell and HP.

The ACCC (Australian Consumer and Competition Commission) is keeping a watchful eye on the merger, though it is unlikely to oppose the deal.

In a further development, Sun has launched a Java app store to offer a range of free and commercial applications developed using the Java language. In its bid to purchase Sun, Oracle said that Java was core to its future business.

***OpenOffice.org* may not be the only project of Sun's to be let go.**

Newsbytes

Google has announced a new product that draws together a range of different communication methods. It's called *Wave* and is in a closed beta at the moment, but Google said most of it would be released as open source eventually.

Google's *Chromium* browser has reached the alpha stage for Linux. The browser is available as a package for Ubuntu as we go to press and features a blisteringly fast JavaScript engine optimised to work with apps such as Gmail.

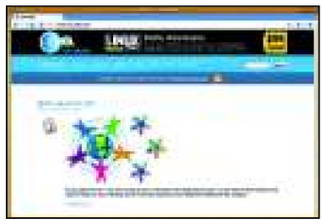

Novell has posted second-quarter profits of over $15 million, up from $5 million for the same quarter last year. This happened despite an 8.5% drop in revenue, which analysts put down to the strong dollar.

The biannual survey figures from the Eclipse Foundation, which manages the release of the Eclipse open source development environment, shows that Linux has increased its presence in both development and deployment of Eclipse applications. Since 2007, the number of enterprises using Linux for Eclipse development has increased from 20% to 27%. Meanwhile, 37% used Linux as a deployment platform in 2007, but this year it had leapt to 42%.

Following a long dispute, Cisco, which owns the Linksys brand, has settled its lawsuit with the Free Software Foundation. It will also appoint a director to ensure that all future developments respect the GNU GPL. The company had been accused of selling wireless routers that included GPL software, but without the sources available as the licence requires. The issue has been ongoing since 2003.

Moblin rebooted

Michael Meeks

This month Moblin 2.0 was finally announced with an incredible new UI for Linux netbooks. This is built around the *Clutter* toolkit, giving flashy 3D effects and a fresh look and feel. We've worked with Intel's team, which did much of the heavy lifting, to polish, optimise and improve the SUSE Linux Enterprise base to be Moblin-compliant and sit under the Moblin UI. Of course, there's an OpenSUSE version of this available from the OpenSUSE Build Service, so you can try out the sleek Moblin UI on your existing system right away, though it does require really good OpenGL support.

It's great to have all this work in public now. The basic technologies come from the Gnome mobile stack. The new panel is built on a *Clutter*-ized *Metacity* window manager, which is also shared with a project to build the next-generation Gnome experience: GnomeShell.

Micro stack

Moblin has a new web browser, predictably built around *Clutter*, which provides a new nagivation experience, with a sweet clustered thumbnail view of your bookmarks and recent pages. Social networking is integrated, enabling you to follow your friends Twittering and easily update your status in the UI. And with the Telepathy-based chat also close to hand, your instant messaging contacts are also only a click away. For email we've written a new, smaller, leaner front-end for *Evolution* called *Anjal*, which provides a simpler, easier-to-use feel suited to netbooks' smaller screens.

We've also created a new media player front-end for *Banshee* called *Cubano*, which takes advantage of Clutter to get fast video playback.

All in all, it adds up to a great pile of innovation that can make Linux the OS of choice on netbooks, fully supported by free software on Intel hardware.

› Michael is a pseudo-engineer, semi-colon lover, Novell *OpenOffice.org* hacker and amateur pundit.

DESKTOP

Mandy in the cloud

Distro offers online backup and storage.

French distro vendor Mandriva is launching a new online service that will seamlessly back up users' documents to the cloud and enable them to be synchronised across various computers. The recently launched service is called *Click'n Backup* and brings cross-platform synchronisation to the Mandriva desktop, since client software is also available for both Windows and OS X. A beta client has been released for other versions of Linux, but it's not yet stable enough for distribution.

In contrast to similar services such as *Dropbox*, Mandriva's offering doesn't have a low-end free option. Users pay €6.99 per month for a 20GB online drive. There are no transfer limits, but the service has an individual file size limit of 5GB. Storage options go up from there, with the top package providing 100GB for €20 per month.

› ***Click'n Backup* features desktop and web clients, so you can access your files wherever you are.**

However, *Click'n Backup* is entering a crowded market. Ubuntu has just revealed its *One* service, while Microsoft has begun to offer 25GB of free space for Windows users and *Dropbox* supports all three OSes. How Mandy's offering fares remains to be seen.

SOFTWARE

Is Linux unviable?

Software experiment struggles to break even.

Developer Bryan Lunduke recently set out to perform an experiment in creating commercial Linux software to see if it would pay his mortgage. He released a trio of applications – an online comic viewer, a comic book editor and an action scheduler – on Linux, Windows and OS X, and then sat back to watch the money pour in.

Unfortunately, things didn't quite work out that way. Lunduke said that after three months his conversion rate among all users – that is people downloading and then paying for software – was 60 to 1. Linux and OS X downloads were broadly similar, but were dwarfed by three times as many Windows downloads. Significantly, although Lunduke offered versions for Ubuntu, OpenSUSE and Fedora, almost 80% of Linux downloads were for Ubuntu, 10% were a non-specific archive and 10% for SUSE.

Lunduke said: "A direct application of the existing shareware model to the Linux marketplace is doable, but I don't think it's overly worthwhile.

› ***Do This Now* is a sophisticated scheduler, but it garnered few sales.**

"This makes it difficult to show a strong business case for larger commercial efforts outside of niches such as virtualisation."

He's not disillusioned, though. "It's possible that this will change over the coming years and it will be attractive and profitable for companies to bring their software to the Linux desktop."

Lunduke's response to the experiment was to change his own business model, making his three applications available free of charge and then offering support, additional paid-for content and soliciting donations to fund further development. He has, for the moment, stopped short of following the open source path, though.

LIMITED PERIOD OFFER

50% OFF

50% OFF great value, feature rich

Web Hosting

Sign up today to receive **50% OFF web hosting for the first 3 months**

Choose **Windows** or **Linux** operating systems and get **unlimited bandwidth** with all hosting packages

HOME PACKAGE

50% OFF

FREE £100 advertising vouchers

- ✔ **UNLIMITED** web traffic (bandwidth)
- ✔ Choice of Windows & Linux OS
- ✔ 2 GB UK web space
- ✔ Over 500 mailboxes
- ✔ MySQL (with Linux)
- ✔ PHP5, Perl, Python, ASP, ASP.NET 3.5
- ✔ Easy website builder
- ✔ Graphical web stats

WAS £4.99PM

£2.49pm

FOR 3 MONTHS

DEVELOPER PACKAGE

50% OFF

FREE £100 advertising vouchers

- ✔ **UNLIMITED** web traffic (bandwidth)
- ✔ Choice of Windows & Linux OS
- ✔ 5GB UK web space
- ✔ Over 1,000 mailboxes
- ✔ 1 MS SQL 2008 database (with Windows)
- ✔ 5 MySQL databases (with Linux)
- ✔ PHP5, Perl, Python, ASP, ASP.NET 3.5
- ✔ SSL secure web space
- ✔ Custom error pages
- ✔ Graphical web stats

WAS £8.99PM

£4.49pm

FOR 3 MONTHS

BUSINESS PACKAGE

50% OFF

FREE £100 advertising vouchers

- ✔ **UNLIMITED** UK web space
- ✔ **UNLIMITED** mailboxes
- ✔ **UNLIMITED** web traffic (bandwidth)
- ✔ Choice of Windows & Linux OS
- ✔ 2 MS SQL databases (with Windows)
- ✔ 10 MySQL databases (with Linux)
- ✔ PHP5, Perl, Python, ASP, ASP.NET 3.5
- ✔ Microsoft Exchange business-class email
- ✔ FREE Outlook 2007 software (R.R.P £99.99)
- ✔ Spam and virus scanning
- ✔ Load balancing – for maximum uptime
- ✔ Search engine optimisation tools

WAS £15.99PM

£7.99pm

FOR 3 MONTHS

SQL Server 2008 · Microsoft .net · MySQL · Microsoft Certified Partner · .net

Unlimited bandwidth with all hosting packages from £2.49pm

50% OFF web hosting for the **first 3 months**

fasthosts.co.uk/hosting

0870 888 3516

24/7 HELPLINE INCLUDED WITH ALL HOSTING PACKAGES

fasthosts

World Class Web Hosting

Prices exclude VAT. Offer correct at time of print.

Comment

Working to rule

Jeremy Allison

Microsoft recently released service pack two (SP2) for *Office 2007*. Normally I wouldn't have noticed, but *Office 2007 SP2* had an important new feature that made me pay attention – it contains Microsoft's first native implementation of the Open Document Format (ODF) created for Sun's *OpenOffice.org*.

And with the implementation of ODF in *SP2*, we finally have one portable office file format, accepted and implemented by most office productivity software. Right? Well, no.

I'm not going to go into great technical detail about what Microsoft did wrong – Rob Weir, chair of the ODF Technical Committee, does an excellent job of that in his blog at **http://tinyurl.com/qedddtd**. In short, Microsoft has actually managed to reduce interoperability between office productivity software by their implementation of ODF inside SP2.

How can this be? Well, one of the reasons is that standards themselves aren't perfect. The ODF standard is missing a proper definition of spreadsheet formulas, which is the truck-sized hole that Microsoft drove through in their implementation. Yet *Microsoft Office SP2* claims to have a fully compliant version of ODF. Although that claim is probably true, as defined by the specification, *Office* is currently useless at interoperating with other products – an attack on the very concept of interoperability.

Trade unions used to have a tactic of protesting called 'working to rule'. It means the punctilious observation of every possible rule in order to disrupt orderly working and it's exactly what Microsoft has done in *SP2*.

I've seen Microsoft do better than this. I've worked with its engineers on CIFS, they've attended interoperability events and even logged *Samba* bugs. Unfortunately, what we see in *Office SP2* is working to rule in every sense of the phrase.

› Jeremy is a lead developer on the *Samba* team and works for Google in San Jose.

LEGAL

Red Hat challenges closed-door MS deal

Swiss government chose MS without bidding process because 'no serious alternatives exist'.

Red Hat and 17 other technology-based companies have issued a legal challenge in response to a large IT contract that was awarded to Microsoft by the Swiss government without a formal bidding process. The consortium of protestors said its challenge raised serious issues about openness and transparency in the Swiss government's handling of public sector contracts and the need for governmental bodies to create a level playing field among competitors.

The disputed contract is a 14 million Swiss Francs deal covering applications, desktops and support for the country's Federal Bureau for Building and Logistics, and is set to last three years.

In response, the government said it didn't need to open the bidding because no serious alternative to Microsoft solutions existed. Citing deployments of Red Hat systems in the Swiss government, the group said it was not acceptable to close the bidding process. Red Hat and its cohorts requested a reversal of the decision and a commitment to hold a public bidding process to enable fair consideration of the merits of open source and other non-Microsoft software products.

Picture: CC Eastside06

› The Swiss Parliament approved a massive IT contract with no bidding process.

Embedded Linux news

› Intel's Atom processor has now found its way into network attached storage. Qnap Systems has launched a Linux and Atom-based NAS system that can run four individual 2.5-inch drives – up to 1.5TB each – on just 18W of power. In addition, the Qnap system will also function as a web server, with packages such as *WordPress* and *Joomla* available to users.

› Asus is expected to follow Qnap's release with the launch of its Eee NAS system. The Eee NAS, a 1TB machine, is expected to be available later this year.

› Qualcomm is set to enter the netbook market this autumn with a range of SmartBook products built around its own Snapdragon chipset. The company said its range would be defined by high-end features such as Wi-Fi, GPS, hi-def video and Bluetooth, but would couple these with a battery life measured in days. The range is expected to cover entry-level netbooks to industrial tablet models.

› A UK startup is taking on the might of Sony and Amazon with a Linux-based e-book reader. The Cool-er is expected to cost less than the Kindle, features a 6-inch, 170dpi greyscale screen and is built on a Samsung ARM400Mhz processor coupled with 128MB RAM and 1GB of storage. Additional space is available via the SD card slot and the device can handle a range of formats such as PDF, EPUB and FB2. It can also play MP3 files.

NETBOOK

Intel debuts Moblin 2

The portable OS seeks to capture the netbook sector.

Intel may be best known for creating billions of CPUs, but it's becoming increasingly worthy of note for its dealings with Linux too. The company has now released an alpha version of its Moblin 2 netbook Linux distro, which is designed to run on Intel Atom-powered netbooks and mobile internet devices (MIDs). Moblin features an untraditional interface that does away with the traditional clutter of panels, menus and desktop icons, creating an overall feel that seems tailored for touchscreens.

While Intel has enjoyed a long-term relationship with Microsoft – and there have been suggestions that it has some significant agreements with Redmond in the netbook space – the company has also forged an alliance with Novell to promote Moblin on the Atom processor. For its part, Novell will build a Moblin-based OS for the netbook sector, which is destined to be sold through its IEM channels, and it will establish a Moblin Development Lab in Taiwan in concert with the Taiwan Institute for Information Industry.

Doug Fisher, vice president of Intel's software and services group, said Novell had been involved in the Moblin programme for over a year, but this new partnership took its involvement further. "The combination of Atom processor-based platforms and Moblin-based Novell software will provide even more opportunities for OEMs, ODMs and the broader Moblin community to deliver excellent mobile internet solutions," he said.

In addition to selling the software, Novell's developers bring a range of expertise and applications to the distribution, including email and media management. Novell CEO Ron Hovsepian said it was important to recognise where the industry was going. "The emergence of such mobile computing platforms as netbooks presents a significant growth opportunity," he said. "We believe that Moblin-based Novell software on Intel platforms will offer OEMs and ODMs exceptional solutions for delivering a full internet experience on such devices."

› Moblin is expected to power a range of netbooks and MIDs, such as the Viliv S5.

Linpus adopts Moblin

Linpus Linux Lite – a distro that ships on a variety of netbooks, including Acer's Aspire range – has just received a Moblin 2-based upgrade to its core product. Linpus was one of the pioneers of the netbook distro, so its jump to Moblin is a significant change.

Linpus said the interface and small footprint of Moblin made it possible for netbooks to boot much faster and that it worked well with smaller screens.

The company demonstrated the new OS at the Computex trade show in Taiwan, but was tight-lipped about when a production version would be released. It said: "The new version of Linpus Linux Lite boots in less than 15 seconds and takes the user experience to the next level: our live desktop gives instant access to recent and favourite websites, communication tools, apps and multimedia."

WEB

SourceForge acquires social network specialist

Ohloh Corp bought to improve software host's community understanding.

The home of many open source software projects, SourceForge, definitively announced it would acquire the Ohloh Corporation in late May this year. Ohloh is a large directory of software and developers, containing information on over 300,000 projects and details of a similar number of individual developers.

SourceForge's group president of media, Jon Sobel, hopes that the acquisition will improve the company's understanding of the open source community and would make it a more effective advertising platform. "We expect the acquisition of Ohloh and the integration of its technology to significantly improve our insights about the open source development community and our ability to target advertising," he said. "The founders and employees of Ohloh are entrepreneurial and well-regarded in the industry. We are excited to welcome the Ohloh team to SourceForge."

Scott Collison, CEO of Ohloh, said SourceForge had been a significant player in the open source development community for years and that an acquisition made sense. He concluded: "We are eager to combine our skills and insights with SourceForge's credibility and reach in the open source ecosystem and look forward to joining the SourceForge team." Financial details about the deal were not disclosed.

› SourceForge is used to manage hundreds of open source projects, and Ohloh should help improve its service.

Comment

The dark art of printing

David Cartwright

Much of what we do in this brave new world of ours is dead hard in theory, but has been made simple thanks to sensible hardware developers and hundreds of software writers. For instance, the IP stack in the average OS is highly complex, yet all we have to do on anything from an iPod to a supercomputer is enter an IP address and we can start a connection. To run a program written in high-level programming languages, such as C or Java, we just type a simple command and a vastly complex compiler takes our source code and translates it (sometimes via intermediate bytecode and a runtime layer) into the ones and zeroes the CPU understands. Why, then, is printing still an art instead of a science?

Don't get me wrong, there's plenty of software out there – not least the excellent CUPS – that makes setting up printers much easier than it might otherwise be. What pains me is that the authors had to write CUPS in the first place to deal with all the wackiness thrown at us by the printer companies. For instance, if I want to print to my HP OfficeJet, there isn't a CUPS driver for that particular model, so I have to pick a driver for another model that I've found is similar enough to work reliably.

We have standards in printer control languages, not least PostScript and PCL. We have standards in printing, too, such as the Internet Printing Protocol (and, for that matter, good old PCL). And there are only so many ways to put ink or toner on to paper. Yet I still have to chant weird incantations at my Linux print configs, and even tweak registry entries on my Windows servers, so they'll stream raw data to a Unix LPD queue.

Surely I can't be the only person wondering why, in 2009, the printer makers can't come up with a simple way of getting our documents down on bits of paper?

› David is an IT consultant with a penchant for cross-platform integration, CTI and proper beer.

DESKTOP

Canonical to expand Ubuntu's app range

New Android software and App Store-like management of apps is coming soon.

Canonical, the company founded by Mark Shuttleworth to marshall the development of Ubuntu, is set to launch two new initiatives that might change the way we acquire and use software on the distro.

Following the trend started by Apple, Blackberry, Nokia and Linspire, Canonical is expected to launch an app store to give users access to a range of free and paid-for software via a unified front-end. Ubuntu has no less than four package management applications in a default installation at the moment, but the new store is likely to be based around the Add/Remove Software option that sits at the base of the main Applications menu. Ubuntu's package management system is already fairly similar to an app store, but the changes are likely to herald better integration of payments and, perhaps more crucially, micropayments into the system. Canonical may potentially also receive a small cut from the sale of commercial software, which would be another logical step on the path to making Ubuntu turn a profit.

The second and more ground-breaking offering is nascent support for apps and applets built on Google's Android operating system on a standard Ubuntu desktop. The Android Execution Environment (AEE) was demonstrated at Canonical's Ubuntu Developer Summit in Barcelona and is thought to be one of the pillars of Ubuntu's expansion plans into the netbook market. The AEE works by convincing the application that it's working within an Android environment rather than on an Ubuntu desktop.

However, Ubuntu's netbook-centric Remix may end up competing directly with Android in this sector. Dell, which ships Ubuntu on some of its netbooks, has recently shown off one of its Mini10 machines running the newly released Android 1.5 OS, which is codenamed Cupcake. LXF

› Michael Frey used Ubuntu's version of *libc* to compile the Android apps.

Linux exams promotion

The Linux Professional Institute has joined forces with Canonical to promote a new discount scheme on the LPIC-1 and Ubuntu Certified Professional exams. The promotion enables participants to sign up for a bundle of courses and save 20%.

Jim Lacey, president and CEO of LPI, said that the course was growing in popularity as Ubuntu becomes more successful. "This is an excellent opportunity for candidates to earn both their LPIC-1 and Ubuntu Certified Professional certifications," he said. "Furthermore, since we now have new LPIC-1 exams, this combined certification status will enable successful candidates to demonstrate that they have in-demand skills and knowledge in today's competitive employment environment."

The courses have been created with the assistance of Canonical's training team. Billy Cina, the company's training program manager, said: "This combined effort demonstrates the strength of the cooperative efforts of open source technology organisations. Together we can advance the cause of open source professionalism while providing real benefits to both of our respective communities."

The LPI has been at the forefront of Linux training since 1999 and has supervised over 200,000 individuals in that time. It operates from testing centres in more than 7,000 locations around the world. To celebrate its 10th anniversary, the LPI will host a range of events at the LinuxTag Expo in Berlin on 24 June. Find out more at **www.ubuntu.com/training**.

DEDICATED SERVERS

FROM £79 PER MONTH

BRAND NEW DELL RANGE WITH WINDOWS OR LINUX

- FREE SAME-DAY SETUP
- 10 MBPS UNMETERED CONNECTION
- NO MINIMUM CONTRACT

	D2-100 £79 per month	D2-310 £119 per month	D2-360 £179 per month	D2-420 £289 per month	D2-840 £499 per month
PROCESSOR	Intel® Pentium	Intel® Xeon	Intel® Xeon	Intel® Xeon	Intel® Xeon
CORES	Dual 2.4Ghz	Quad 2.4Ghz	Quad 2.66Ghz	Quad 2.83Ghz	2 x Quad 2.5Ghz
MEMORY	1GB ECC	2GB ECC	4GB ECC	8GB ECC	16GB ECC
HARD DISKS	1 x 160GB SATA	2 x 250GB SATA	2 x 500GB SATA	2 x 1TB SATA	2 x 450GB SAS
RAID	None	Hardware RAID	Hardware RAID	Hardware RAID	Hardware RAID
CONNECTION	10 Mbps	10 Mbps	10 Mbps	10 Mbps	10 Mbps

OR BUILD YOUR OWN FULLY CUSTOMISED DELL DEDICATED SERVER AT: **www.redstation.com**

- DELL POWEREDGE SERVERS
- WINDOWS OR LINUX
- FREE SAME-DAY SETUP
- 10 MBPS UNMETERED
- NO MINIMUM CONTRACT
- REMOTE SERVER CONTROL
- SECURE PRIVATE NETWORK
- 24/7 TELEPHONE SUPPORT
- PRIVATE UK DATA CENTRE

CALL FREE ON: **0800 622 6655** OR VISIT: **www.redstation.com**

DELL REGISTERED BUSINESS PARTNER

redstation

Yahoo Hack Day 2009

Code, caffeine and beanbags aplenty at Yahoo's 24-hour programming fest...

What happens when you pull hackers away from the internet and get them together in real life? They hack, of course! On 9 May at the Congress Centre in London, Yahoo held its annual Hack Day event, an opportunity for developers from all walks of life to get together and work on the Next Big Thing. The event was "deliberately unstructured" – developers work independently or in groups, and have the freedom to work on what they want, within the confines of three rules:

1. The project must be based on an existing Yahoo API or project.
2. Completed hacks must be submitted for demonstration by 1.30pm the following day.
3. Developers must show a 90-second demo of their work.

Hack Days are based on events that Yahoo used to hold internally, wherein employees could share ideas and code without being restricted by their regular schedule. In recent years Yahoo has held public ('open') Hack Days in Taiwan, India and Brazil, so when it came to London we wanted to hop over and see what it was all about.

> **"It'll have millions of users – the next Twitter."**
> **Hacker Dom Hodgson on GeoPubbing**

24hr party people

Coders gathered around tables in the main room in the morning, grabbed a quick spot of lunch, then started a 24-hour marathon hacking fest at 1.30pm. And really, it's 24 hours – Yahoo packed the room with bean bags and mats for weary-eyed developers to grab power naps.

It wasn't all hands-on hacking though, with separate sessions on the state of web technologies. Rasmus Lerdorf, creator of PHP, gave a talk on using the language with XML, JSON, YQL, YUI and other TLAs, while we were intrigued by *Mozilla Bespin*, a web-based code editor whose name "officially has nothing do do with Cloud City in *Star Wars*" according to Joe Walker, one of the four developers.

Although *Bespin* is still in the early stages of development, it handles large files with ease (Joe demonstrated a 26,000-line file scrolling around smoothly in the editor within *Firefox*) and has some big plans for collaboration. After the demonstration we sat down with Joe to find out more about the project – we've put the interview online at **http://tinyurl.com/trbespin**.

Show me the code

Back in the coding zone, things got off to a shaky start thanks to the lack of wireless internet access. For some developers working on desktop apps

A purple double-decker bus, yesterday.

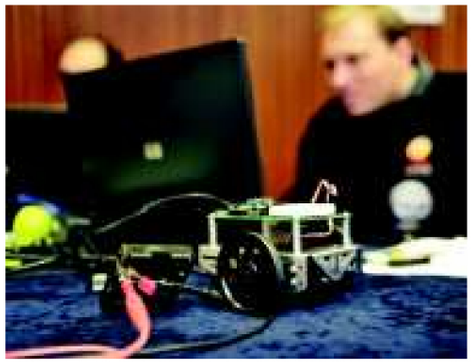

The many robots in the hall proved that Hack Day isn't just about the web.

› "Hands up if you've not slept in 23 hours and your eyes have gone funny."

(the majority of laptops we saw were Macs) that wasn't a big problem, but other teams just had to dream up grand ideas until the network packets started flowing after a couple of hours.

Dale Lane was piecing together an idea called Guest Pass. This started as a scribbled diagram that he explained to us. "It gives people access to location info for a limited time, and they don't need to sign up to anything." We tried this out with Yahoo's Fire Eagle service, which lets people share their location information but currently requires a Yahoo account to use. With Guest Pass you can give someone one-off access to your location via Fire Eagle, and they don't need to be signed up. **http://feguestpass.appspot.com**.

We instantly fell in love with the work of Dom Hodgson and his team's GeoPubbing project. As you might expect, this combines geocaching and pub crawling using Fire Eagle. You get a geocode for a pub, find your way to the pub, have a drink then somewhere in the pub you'll find a QR code which contains the location of the next pub to visit. According to Dan: "it'll have millions of users – the next Twitter!" **www.newmedias.co.uk/geopubcrawl**.

Happy/sad

Two chaps who identified themselves only as "Squirrel and Chris" were scheming with some Lego Mindstorms NXT kit. When we saw them, they were just starting to assemble smiling and frowning robots with grand plans to hook them up to the internet. One goal was to "get two Twitter accounts and make them argue" – because there's nothing like the internet to get arguments going! There were more mechanical shenanigans on another table, where a bunch of Belgians were dreaming up "a robot to get beer from the bar". It wasn't clear how the Yahoo APIs came into play in this...

The four winners were announced on Sunday. IntelliSearch (**http://fluttercookie.zandrock.com**) melds a Dasher-like accessibility keyboard with Yahoo's predictive search technology, while They Work For EU (**http://theyworkfor.eu**) aims to provide translations of proceedings in the European Parliament. Boss of MySpace (**http://tinyurl.com/pq2wb4**) uses Yahoo's Boss (Build your Own Search Service) to find music on MySpace and snag tunes without going to the page, and OpenFreecycle (**http://dharmafly.com/openfreecycle**) lets you search for stuff that people are giving away via the planet-friendly Freecycle system. LXF

Next Month

Don't miss next issue for our full interview with PHP creator Rasmus Lerdorf!

United Linux!

Meet up with like-minded open source fans – both locally and internationally.

Not enough spice

After reading our *Spice Up Your LUG* feature online, the mysterious Morten Juhl-Johansen Zölde-Fejér pointed us in the direction of an article they'd authored on Technographer.net earlier in the year. It's an article entitled 'Arranging Linux Events', and it's a great read if you're looking for more ideas on getting a group started. It starts with the basics – finding out potential numbers, the pros and cons of registering interest and finding a venue. But it quickly raises an interesting point we hadn't mentioned in our original article, and that's that you need to ensure any venue you choose has enough power options for your group. If not, you've got to supply your own.

Publicity

Another point we didn't discuss in the original article was advertising your meeting, and while your LUG's mailing list is the first obvious choice, it's not going to grab new members very easily. Morten suggests broadening your group's appeal, perhaps including local BSD enthusiasts, or even folks from the burgeoning mod/DIY tech community. As long as everyone shares a common interest, this diversity is likely to make a meeting feel more vibrant, and there will be more breadth in the topics that can be discussed. To close the article, Morten dropped a great quote attributed to an anthropologist called Gregers Petersen, "The difference between this happening or not happening is the difference between you doing it or not." And we couldn't agree more

graham.morrison@futurenet.com

Healthy FOSS

Houston doesn't have a problem with open source.

Fred Trotter dropped us a line to tell us about an open source event being held in Houston, Texas. It's called the 'FOSS in Healthcare Unconference', and rather than being driven by a strict schedule and expensive invited speakers, the format of an 'unconference' is typically governed by the attendees. The environment is a far less formal setting than the average conference centre too. The catalyst for this particular gathering is the specific mention of open source in the hi-tech portion of the US Government's Stimulus Bill, an economy boosting initiative recently signed into US law by the President Barack Obama.

The conference will feature the top open source systems in healthcare, and aims to promote its use and feasibility in an area that needs some help. Highlights include *OpenMRS*, a community-developed electronic medical record system, and *Mirth*, an open source interface engine that has had great success replacing entrenched and proprietary software used by great swathes of healthcare installations. The event is hosted by the Free Medical Software Foundation, and sponsored by the HealthQuilt project. It's being held at the HAL-PC centre in Houston at the end of July/early August, and ticket and accommodation details can be found on the project's website. Drop us a line if you're attending and we'll happily print a report on how the event went. We just wish there were something similar in the UK that NHS IT management could attend before wasting £20bn.

http://fosshealth.eventbrite.com

Check the health of open source at the end of July in Houston, Texas.

Community news

It's conference season!

Now that the summer is upon those of us in the northern hemisphere, it's time for some conference action. In the US, O'Reilly's Open Source Convention is taking place as you read this, on 20–24 July. After six years in Portland, this year's event is being held in San Jose. During those six years, attendance grew from 1,300 registered attendees in 2003, to 3,000 registered attendees in 2008, and even Tim O'Reilly was turned away from a couple of sessions due to of the lack of space. The new venue should give OSCON enough room to grow over the next few years, as well as a change of scenery. We're hoping to attend this year's conference. so keep an eye on **TuxRadar.com** and these pages for our reports and experiences.

In what is perhaps a sign of the growing free software movement, the Linux World Expo has changed its name to 'OpenSource World', although it's staying in the same location this year – the Moscone Centre in San Francisco.

While OSCON covers the geekier side of open source software development and emerging technology, OpenSource World has a greater concentration of businesses. But they're both great events if you happen to be in the Bay Area of California and can afford the entry prices.

en.oreilly.com/oscon2009

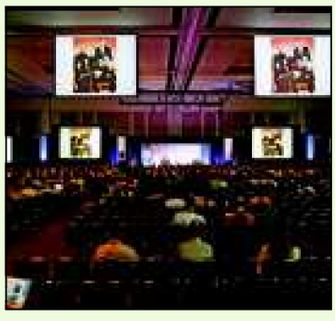

We're heading off to San Jose this year to experience the new OSCON location for ourselves.

Mailserver

Write to Paul at *Linux Format*, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or **lxf.letters@futurenet.co.uk**

Always the sun

There's a new release of OpenSolaris coming soon. Could we have it on the cover disc?
Rob Jones

Mike says: Oh go on then, because you asked...

Debian doo-doo

When I saw that LXF118's DVD offered the latest Debian distro, I wondered if I might finally pension off my old Xandros 4 Open Circulation.

Oh dear. The graphical installer was a mess. The English in the help was often obscure or ambiguous, the buttons marked Go Back took me onwards and the guided partitioning suggested unhelpful options that ignored my existing Windows disk. But the real let-down was that after two attempts, I hit a brick wall: no sign of a functioning operating system and no hint of what might have gone wrong. All I got was a simple prompt.

What a relief to go back and reinstall the underpraised Xandros! Everything went smoothly, the guidance was crystal clear and I wasn't constantly asked to make pointless 'choices' *à la* Debian.

Obviously many people will be able to install Lenny without difficulty, but how many more will give up when confronted with a distro that takes them back to the disappointments of ancient products such as the unlamented Mandrake ?
Maurice George

Paul says: Wow, that's a pretty negative experience, and I'm surprised that Debian fared so poorly given how much testing it received. Please do take the time to submit a bug report or two, because it will help many others!

Letter of the month!

Catch up

Wins T-shirt

Since I first started reading it from LXF100 I've been a huge fan of your magazine. It was your publication that converted me and (indirectly) my girlfriend, mother, grandparents and best friend into fully fledged home Linux users. Since then, I've become a subscriber and constantly try out new distros from your cover discs.

I noticed recently that Channel 4 has updated its on-demand service (even though they advertise it as now available on your Mac) it is now compatible on a Linux desktop by using Flash just like the BBC iPlayer. I feel quite proud by this fact, as I kept pestering them with emails asking "When will your service be cross platform" and being met with the reply "We're going to get around to it".

One more thing – I was browsing my local recycling centre the other day and to my pleasant surprise while looking through their collection of DVDs and computer software discs I discovered this small collection of vintage *Linux Format* discs, all going for £2 each.
Stephen K Fuller, Scottish Highlands

Unreal Tournament, Beneath a Steel Sky, OpenOffice.org 1.0...

Paul says: Welcome aboard! Although Mike is convinced that £2 is clearly too cheap for LXFDVDs, we're just glad to see free software being spread around a little bit more. And well done for writing in asking for Linux support – if everyone expressed their views like this, the software world would be a very different place indeed!

Music to your ears

I wrote this while listening to the latest TuxRadar podcast (Episode 9). I just want to say that these are amazing! Keep up the good work.
Sam Tarling

Paul says: It's great to hear you're enjoying our podcasts – we release them every two weeks, and hope to keep them fresh with new ideas all the time. Let us know if you have any suggestions!

TuxRadar

Tree of life

Your article in LXF119 about Slackware and its first release in 1993 stirred an emotion inside me. Whatever happened to Yggdrasil? I know, I know – I should go take a look on Wikipedia and find out for myself. But hey, I'm lazy, and I know that now I have mentioned it I will have hit a nerve deep inside your grey matter and you will be wanting to reply.

Seriously though, I still have my original Fall 95 Edition, two-CD version with the manual, and I'd be interested to know how many readers of LXF remember Yggdrasil Linux.

Sadly though I never, ever, got it working correctly at all on either my 386 DX40 or my 486 SX25, both with Cirrus Logic Truecolor video cards. It would not allow the correct graphics format and insisted on MCGA mode, (mode 13H). I would have settled for VGA mode, (mode 12H), as in Windows 3.xx but alas it was never to happen. Its screen looked awful, much like Windows 3.xx and would not scroll like the Amiga did, post Workbench 1.3x days, so part of whatever was on screen was always hidden.

It was totally unuseable and was my first taste of Linux, which was sour, especially when Windows 95 easily installed on the same machines and looked much better... How things change eh?
Bazza

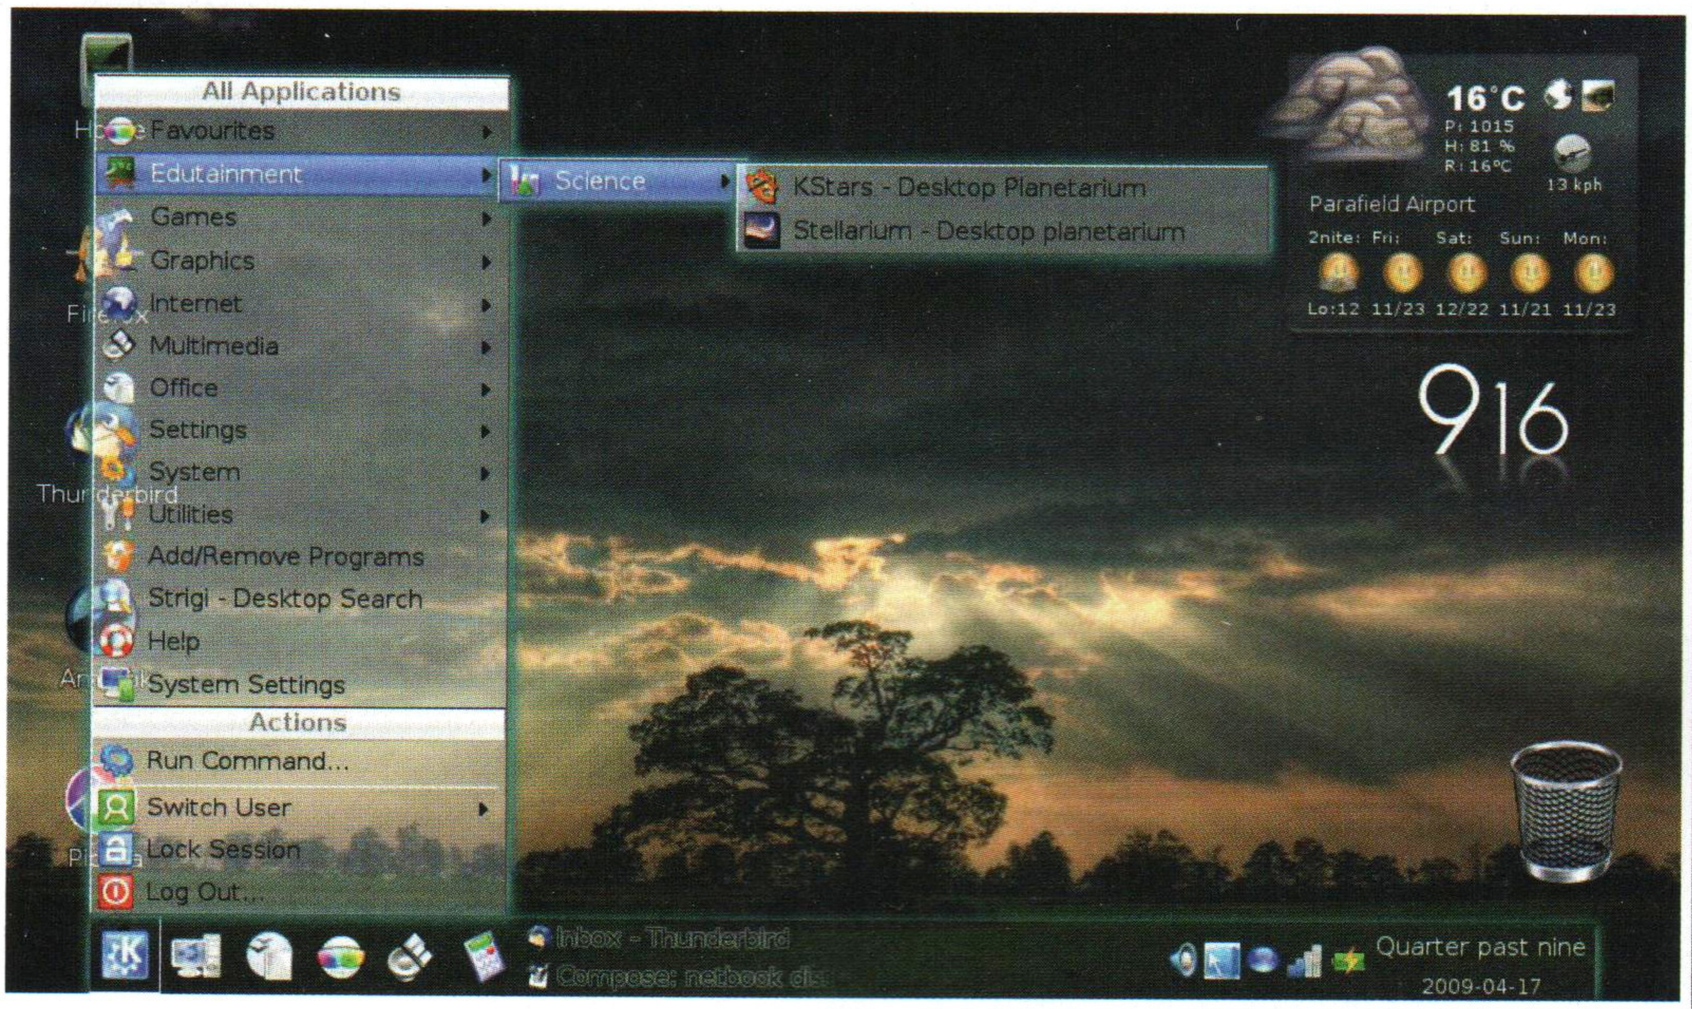

› **After the disappointment of early netbook interfaces, Peter's Kubuntu setup shines like a diamond.**

Paul says: OK, veteran readers: this is your chance to reminisce. Send in your "I first used Linux..." stories and we'll print the best!

Shiny things

Behold my Acer Aspire One (8GB SSD) running Kubuntu 8.04.1 (above)! Full *Compiz Fusion* desktop effects running sweet on 1.5GB RAM. During install I chose an 8GB SDHC card as my home partition, so root has the whole SSD to itself. It runs sweet, looks cool and collected – and my wife uses it every day, I can't get near it! Bugger...

Peter Henderson, Hillbank, South Australia

Paul says: Fortunately netbooks are cheap enough that you can just buy another one!

Join us

Inspired by getting the letter of the month in LXF114, I have created a MySpace profile using only Linux and a 1.3-megapixel cameraphone.

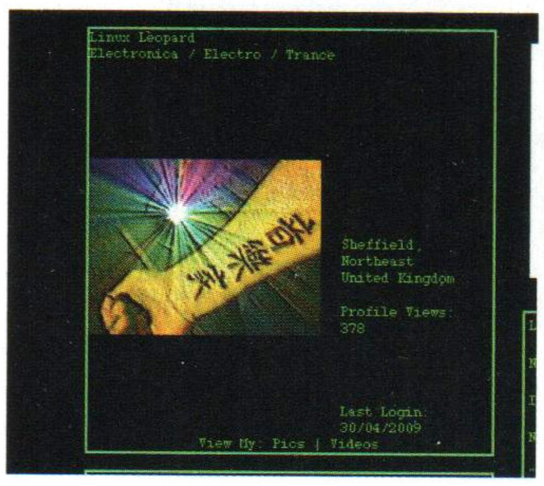

› **Do people still use MySpace? We're all on TwitBook these days.**

As I've seen at least one shameless plug for a site in Mailserver, I would ask you, please, to mention my site **www.myspace.com/linuxleopard**.

It has been and on-and-off project for six months and I would welcome any Linux friends (with or without comments) on my profile.

It may be worth mentioning that getting *Adobe Flash Player* to work on Ubuntu simply requires an install of the appropriate wrapper using *Synaptic*. For anyone who remembers DSSI Fun there's a remix on the website.

Paul

Paul says: We've done our part and printed your URL. If you're looking for MySpace friends, readers, sign up above!

Remote control

I continued my subscription with *Linux Format* when I moved out to New Zealand six years ago and look forward to every issue. I have just been enjoying LXF119 and in particular the excellent article on SSH/VNC and remote working. If I could correct your references regarding NX, there are two versions – NX from NoMachine, which has a free, two-user offering as well as the regular paid-for version, and a separate (but based on same technology) open source version called FreeNX (note the confusing difference – Free NX and FreeNX) which is not limited to only two users, but lacks the sophistication and bells and whistles of the proprietary version. How do I know? I'm returning to the UK for a holiday later this year and wanted to ensure that I could completely access my home server here from my Ubuntu NBR netbook and after a lot of research and even more setting up, I can now achieve that, with a high level of security via SSH and much better performance via NX technology than you can get with VNC.

On another topic, while I appreciate the importance of personal choice, I think that one of Linux's Achilles heels is exactly that – multiple inconsistent ways of providing maintenance (*Yum*, *APT*, *RPM* etc) and multiple desktop platforms discourage both developers and newbie users from taking the leap to a better environment. Despite not being quite perfect yet (whatever that is) the Ubuntu offerings seem to have made the most progress towards consistency and easy of use (which, for all its sins, is how Windows got where it is today).

Martin Sydenham (ex-IBM training manager for IBM Linux offerings to resellers, and developers across Europe)

Paul says: Thanks for the correction! As for packaging, you're quite right, and things don't seem to be improving!

Up the worker!

First, I wanted to add my own favourite to the orthodox file mangers list [from LXF119's Roundup], one that I believe meets or exceeds the published criteria for inclusion. The program is called *Worker* and is available in (GPL) source and a number of binary flavours including a port for Mac OS X) at http://boomerangsworld.de/cms/worker/download?lang=en.

***Worker*'s highlights include being pure *X Window System*-based (no Gnome or KDE required) and thus is extremely light weight and portable. It has built-in scripting hooks that can fully integrate with shell programs, so there is little limit on what criteria can be used for file list display, selections and operations. Any path or script or built-in function can be assigned to the fully configurable panels of menu buttons for ready access. Also of particular note is *Worker*'s support for the AVFS virtual filesystem for transparent network support. *Worker* source package compiles and installs easily with make and make install and has very few dependencies. And re: Over To You, yes! I would love to dump Konqued-out and NOTaius for a more practical, compact and functional OFM! With a just little scripting, I think *Worker* could maybe meet the last couple of criteria of direct data DVD burning and URL browsing, too!**
Aaron Ruscetta

Paul says: *Worker* seems like a very capable file manager-cum-desktop environment, and should prove welcome to folks who are concerned with memory usage.

Life in the old dog

I have been running a Lenovo 3000/C200 laptop as my primary computer for almost three years. It is 100% Linux with Ubuntu 8.04 as the OS. I also keep an old (last century) IBM Aptiva which I use to check out live CDs and to try to install various small distros. Just for fun the other day I decided to install the latest PCLinuxOS, not expecting to have any success. Imagine my surprise when the distro loaded without any hiccups. Sure it's slow – the Aptiva only has a Celeron 666MHz – but it runs! It even detected that the in-built SIS graphics would be better than the Tseng add-on and runs it at a crisp 1024x768. I tried Xubuntu on this box and it could only handle 800x600. I could easily switch to PCLinuxOS as my primary OS and may do so if and when I upgrade my laptop
Tony Moloney, Western Australia
PS Here's a photo of me, my wife and Tux as we boarded the SuperStar Gemini on its last trip out of Australia.

The kids are alright

In the print equivalent to Trevor McDonald's 'And finally', we look to the Linuxers of tomorrow...

Child's play

Thank you – I took your advice about a kids' distro in lxf118 and installed Quimo on two old computers I had for my grandchildren (ages five and eight). Previously I had Edubuntu, which turned them off straight away but with Quimo I can't get them off the computers. Their main problem with Edubuntu was that it wasn't Windows like at school and they couldn't get their heads around the new system, but with Quimo's easy interface they've just dived straight in and used it. So two new converts to Linux looks on the cards.
Albythom, Tyneside

Come on tiger!

A pic of my little boy getting his hands on a new issue before me! Doh, I know I shouldn't have gone to work that day. Now that's got to be worth a T-shirt?
Dan, Poole, Dorset

PS: His fave distro is Ubuntu 8.10

Paul says: Thousands of Mac fans around the world have long wondered about the identity of Mr Blurrycam, the super spy who takes sneaky photos of unreleased Apple hardware. And now you have revealed yourself to LXF, Dan, so expect a deluge of requests from Mac fanboys! As for Qimo, we agree it's a great respin, and deserves more love from us.

All of us were this delighted by the release of Debian 5.0.

Helpdex

shane_collinge@yahoo.com

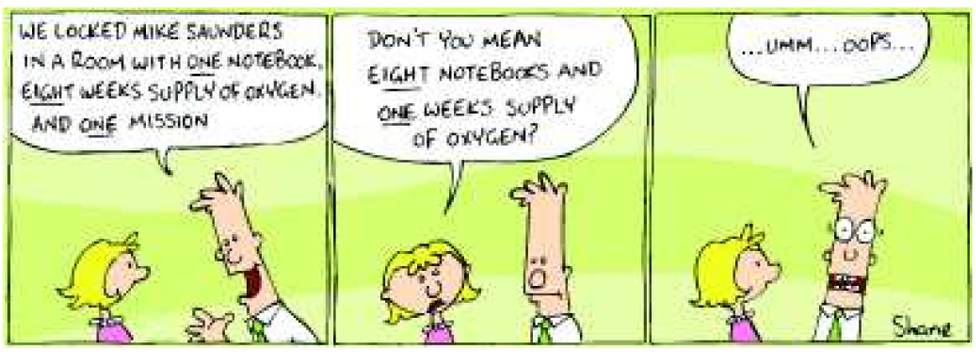

Paul says: Good choice! PCLinuxOS has a knack of getting hardware support just right.

Amaz oh no!

Kevin Atkins' letter on Amazon and downloads is unintentionally misleading. I am assuming he is using Ubuntu, or SUSE, or maybe Fedora. Well, I use Mandriva, and I cannot download MP3 albums from Amazon. I contacted Amazon about this, but never really got a proper reply. In fact the first reply wrongly told me that any Linux system would do. The second one simply stated which systems I had to use. I just cannot understand why they should limit their customer base in this way. This is not a case of me wanting to avoid proprietary codecs – I do not really care how I download the stuff. I just cannot do it.

But Amazon's email to me did state: We're Building Earth's Most Customer-Centric Company. Well, they have a long road ahead!
Jorgen Stepputat

Paul says: I don't think we can blame Amazon – if something works with Linuxes A, B and C, but not D, E and F, that's Linux's fault!

UK broadband

I've read on the BBC that the UK government is thinking about guaranteeing a minimum broadband speed of 2Mbit by 2012 across the country. I'm a bit curious about this, as the standard speed in Sweden was set to 5Mbit a few years back, and I haven't been able to find anyone not offering 8Mbit the last two years. I'm just curious as to what is standard in Europe with regard to internet speeds?

In Sweden one company is now building a 150Mbit mobile network, and the others will probably buy in or invest in it.

Since I found that Linux works essentially only when in a network – you can work without from DVD installs but how often do get everything you want from them – the future of Linux depends on a fast net to spread and get easier.
Werner Johansson, Sweden

Paul says: Most of the UK gets fast broadband, but the problem lies in the more geographically challenged parts of the country that are cut off from current services – these are the places that will benefit from the move.

What's in the box?

With trembling fingers I opened LXF118. All about Debian. Cooool! Everything I always wanted to know, like why it's so huge – 31 CD ISO images at about 650MB each in the repository. More than 20GB. Holy smoke, how do you decide what to use? No worries, 13 exclusive pages in LXF119 spread over eight articles with contributions by nine super-geek heros, Steve McIntyre and all. But alas, having pored over each word I still don't know what is in the 30 CDs. Does anybody know?

This speculation raised a further question: If Debian is the mother of so many popular distros, including ubiquitous Ubuntu, why don't experts (like you fellows) recommend more/most serious Linux users to just use Debian. I tried Steve's installation and I suspect it is a better product than Ubuntu because it already has a slew of applications that Ubuntu does not. Or are those just some nice extras courtesy of Steve? Notwithstanding this paucity of solid enlightenment you produce a superb Linux mag and every month the fingers will tremble.
Mike Ohlson de Fine, Zululand, South Africa

Paul says: We recommend Debian whenever it's appropriate – which is usually when long-term reliability is key!

Men, women and penguins can co-exist peacefully. But not with fish.

debian

Is Debian still relevant, or has Ubuntu stolen its thunder?

Desktop Linux

I read Susan Linton's latest Distrowatch column with interest and find myself agreeing with her. I suspect that 'the year of desktop Linux' means different things to different people. I've been using Linux as my main desktop OS for years now and it does everything I want it to. I surf the web, do my banking and my shopping, share files, download stuff, email, watch films, organise and listen to music and radio, organise and edit my family photo album, and so on.

That is more than enough for myself and, I would guess, most other desktop users too. When will the year of Desktop Linux arrive? For me, it arrived about three years ago.
Huw Taylor, Rushden, Northants

Paul says: You're right, Huw, and I think the world is starting to agree with you and Susan.

Play nice

As a gamer I was interested to read your review of *CrossOver Games*, as anything that enables me to finally ditch that Vista partition would be received very greatfully. You gave the product an 8/10, but I'm still none the wiser on whether to part with my money or not.

Sure, it runs supported games well, and unsupported games with varying degrees of success, but so do *Wine* and *Cedega*. Without any kind of comparison to the alternatives, it's very unclear as to whether it's worthwhile. I'd really like to know what the benefits are.
Simon

Paul says: OK, we'll try to set up some benchmarks for you! LXF

Ultra-fast dedicated servers with Intel® Core™ 2 Duo processors

With a choice of operating systems and unlimited bandwidth, Fasthosts dedicated servers are designed to deliver the performance you demand.

Save now with 50% OFF our most powerful servers for 3 months!

DEDICATED SERVERS FROM £34.50 ~~£69~~

	DS 250	DS 450	DS 650	DS 850
Processor speed	2 x 1.8 GHz	2 x 2.13 GHz	2 x 2.4 GHz	2 x 2.66 GHz
RAM	512 MB	1 GB	2 GB	3 GB
Hard drives	80 GB	250 GB	2 x 250 GB	2 x 500 GB
Bandwidth	UNLIMITED	UNLIMITED	UNLIMITED	UNLIMITED
Special Offer		SAVE 50% FOR 3 MONTHS		
	FROM £49pm	FROM ~~£69~~ £34.50pm	FROM ~~£99~~ £49.50pm	FROM ~~£189~~ £94.50pm

- ✔ Independent remote management
- ✔ Root access – total administrator control
- ✔ Choice of Windows & Linux OS
- ✔ Unlimited bandwidth
- ✔ MySQL and New MS SQL 2008 databases
- ✔ Your server up and running in minutes
- ✔ FREE SETUP when you pay annually
- ✔ UK based data centres

24/7 UK phone & email support

Microsoft SQL Server 2008 NEW

Save 50% for 3 months • Unlimited Bandwidth • 24/7 UK phone & email support

Get your server online in minutes

fasthosts.co.uk/dedicated

0870 888 3505

Prices exclude VAT. Offers correct at time of print.

fasthosts

World Class Web Hosting

LINUX FORMAT Reviews

All the latest software and hardware reviewed and rated by our experts

GRAHAM MORRISON no longer needs a whole canister of Silvikrin hairspray to look like this.

So update me

One trend I really dislike, and one that shows no signs of abating, is the tendency for software and hardware developers to release now and fix later. Many games consoles, televisions, printers and phones are sold containing a slew of bugs, with the manufacturer safe in the knowledge that when complaints reach a critical mass, problems can be patched. This is a poor excuse for a lack of proper testing. Or any testing at all, if my television is anything to go by.

In the good old days before the internet provided a bailout clause, the developer had no other option than to ensure their product worked as advertised. What's worse is that in many hardware reviews I've read recently, the reviewer will often excuse the lack of a certain feature because there's some vague promise of its addition in the future.

KBugFix

Which is why I find the decision taken by the *KOffice* developers to release version 2.0, despite a huge caveat that it's "not aimed at end users", so bizarre. Have they learned nothing from the painful transition to KDE 4? I have no problem with plenty of early releases for testing, but I do feel there needs to be a clear demarcation of quality when it comes to major releases. How else are unsuspecting users supposed to differentiate between a work in progress and the final version? I'm now looking forward to *KOffice 3* being released next week. Suitable only for dreamers.

graham.morrison@futurenet.com

Our pick of this month's releases:

GP2X Wiz..... 24
The return of the open games console that can run anything from console emulators to *XBill 2.1*. But can it improve on the woeful score we gave its predecessor?

Kdenlive 0.7.4 26
Suffering from the opposite condition that afflicts *KOffice* developers, this insignificant version number marks a major release of this video editing suite.

Ulteo OVD..... 28
If your friends can't commit to a full Linux install, use the latest release from Gaël Duval's team to give them a Linux desktop through a web browser.

Jets'n'Guns..... 29
Those of you who think the best shooting games were *R-Type*, *Nemesis* and *Side Arms* can rejoice. The sideways scroller has made a comeback thanks to LGP.

EnergyXT 2.5 30
Possibly the best integrated electronic music composition tool on the Linux desktop. Stuffed full of new features and drag-and-drop effects, synths and sounds.

OpenOffice.org 3.1 32
A hare one day ridiculed the short feet and slow pace of the tortoise, who replied, laughing: "Though you be swift as the wind, I will beat you in a race." Or will it?

Book reviews p34

GP2X Wiz p24

The dodgy joystick control of the GP2X has been replace by a direction pad.

Jets'n'Guns p29

Prepare yourself. When the LHC finally gets repaired, we may lose a dimension.

Our verdict explained

All the products that we review are rated from 0 to 10, with 10 being highest. The categories we use to rate products are usually features, performance, ease of use and value for money, but software available without a charge might be rated on documentation instead of value for money. Regardless of the four criteria used, we always give an overall score out of ten. Products that stand out from the crowd may receive our prestigious Top Stuff award. Only the best is considered for this award – scoring highly isn't enough by itself.

When reviewing free software, we will usually use the recommended distribution of the software. In some cases this will be hand compilation using *GCC*, but if the developers recommend Autopackages these will be used instead.

LINUX FORMAT Verdict

Google Earth

Developer: Google
Web: http://earth.google.com
Price: Free under proprietary licence

Features	10/10
Performance	9/10
Ease of use	9/10
Documentation	9/10

» *If all the world's a stage*, Google Earth *is the theatre. Easy to use, utterly addictive and reassuringly practical.*

Rating 9/10

GP2X Wiz

Simon Pickstock goes misty-eyed with nostalgia, with this super-geeky console and arcade emulation platform.

In brief...

» Linux-based open source portable games platform for retrostalgic, gaming delight.

Specification

- **CPU** 533MHz
- **Screen** OLED
- **RAM** 64MB
- **Storage** 1GB
- **Supported formats** MPEG4, DIVX, XVID, JPEG, BMP, GIF, PNG, WAV, OGG, TXT, Flash 8.0.

Everyone loves a bit of retro, and the GP2X Wiz serves up lovely great dollops of it. An evolution of the GP2, the Wiz features a touchscreen with a stylus, an ARM CPU running at 533MHz and a memory card slot for game storage. Most importantly of all, it's based on Linux. The Wiz connects to your PC using a proprietary USB cable, which also acts as a charger, and when connected, asks if you would like to use the built-in memory or SSD card for transfer. Your chosen destination will mount on the desktop as a removable storage drive, enabling you to copy files across.

The Wiz's OLED screen is nice and bright, and the interface consists of just six icons; SD card, built-in games, flash games, entertainment, launcher and settings. The built-in games are simple clones of old classics, such as *Snake*, *Tower Defence* and so on. There are five flash games, but as we don't speak Korean we can't tell you what they're called. Still, this is pretty irrelevant, as each of the games is of the basic brain-training type, where you have to match symbols, or perform simple arithmetic. The entertainment section enables you to play audio, video and picture files.

> **"At a primeval geek level, the Wiz pushes all the right buttons."**

The games require the use of the touchscreen and stylus, although we found that the touchscreen was simply not sensitive enough, requiring several taps to register. The net result was that our brain age was given as 61 – gah!

› **The GP2X Wiz fits into a baggy trouser pocket – just take your keys out before you sit down.**

› **There's nothing like a trip to Spring Yard Zone to unleash your inner child, but please stay legal! Did we mention this was written from our moon bunker?**

Further investigation revealed that the mapping of the touchscreen was at fault, and once this was recalibrated, things ran a lot more smoothly, and we even knocked 20 years off our brain score – get in!

Nostalgic Shangri-la

Several console emulators have already been ported to the Wiz, including the Megadrive (Genesis) and SNES, as well as *Mame* and more will no doubt be ported soon. However, in order to find them, you'll have to do some searching around, as the UK site for the Wiz currently has no working links, and most of the content is to be found on forums that require registration or is held on fan sites.

Thanks to Mike's stash of ROMS were soon whiling away the hours with *Sonic the Hedgehog* and Road Rash, using the Genesis emulator. Ahh, bliss. However, this did bring to light a rather serious design flaw. While the controller on the left is a traditional four-way controller, the right-hand one is four separate buttons, arranged in a four-way pattern. The buttons are tiny, and placed far too close together for comfort. In addition, every time a new level loads, the volume resets to maximum level, which is annoying if you're in public and downright painful if you're using headphones.

It's also worth pointing out that it's an exercise in frustration to find the emulators, copy them across, get them running then find ROMs you can load. While the Wiz is clearly meant to run old console games, the ROMS are of course copyrighted, so finding them is a challenge in itself, not to mention legally dubious. In the end, our trip down memory lane, was more akin to an attempt on Everest's north face than the walk in the park we were looking for.

At a deep and primeval geek level, the Wiz pushes all the right buttons (albeit badly laid-out ones), and if you want total control over your emulator gaming without having to flash your mainstream handheld games console, this is the ultimate toy. On the other hand, if you're a more casual gamer and don't want to spend half your life searching the intertubes for dodgy content, stick with your PSP or DS. LXF

LINUX FORMAT Verdict

GP2X Wiz

Developer: Green Park Holdings
Web: www.globalgph.com
Price: £130

Features	8/10
Performance	8/10
Ease of use	7/10
Documentation	8/10

» *Requires a large investment of time to deliver the best results.*

Rating 7/10

HOSTING NEXT LEVEL

NEW!

HETZNER DEDICATED ROOT SERVER FROM

£39 PER MONTH

SAVE £7.50 AS A NEW CUSTOMER!
Using this voucher code in your order:
251206 (Expires 09 July 2009)

HETZNER DEDICATED ROOT SERVER EQ 4

- Intel®Core™ i7-920 Quadcore incl. Hyper-Threading Technology
- 8 GB DDR3 RAM
- 2 x 750 GB SATA-II HDD (Software-RAID 1)
- Var. operating systems
- Unlimited Traffic*
- Rescue system
- Image installer
- 100 GB Backup space
- Reset service
- No minimum contract
- Setup Fee £119

£39 per month

HETZNER DEDICATED ROOT SERVER EQ 6

- Intel®Core™ i7-920 Quadcore incl. Hyper-Threading Technology
- 12 GB DDR3 RAM
- 2 x 1500 GB SATA-II HDD (Software-RAID 1)
- Var. operating systems
- Unlimited Traffic*
- Rescue system
- Image installer
- 100 GB Backup space
- Reset service
- No minimum contract
- Setup Fee £119

£59 per month

HETZNER DEDICATED ROOT SERVER EQ 9

- Intel®Core™ i7-965 Quadcore incl. Hyper-Threading Technology
- 12 GB DDR3 RAM
- 3 x 1500 GB SATA-II HDD (Software-RAID 5)
- Var. operating systems
- Unlimited Traffic*
- Rescue system
- Image installer
- 100 GB Backup space
- Reset service
- No minimum contract
- Setup Fee £119

£79 per month

* Traffic usage is free. We will restrict the connection speed to 10 MBit/s if 2000 GB/month is exceeded. Optional, a permanent 100 MBit/s bandwidth will be charged at £22 per additional TB.

HETZNER ONLINE

Hosting Next Level means that Hetzner Online now offers the most powerful dedicated hosting solutions currently available. The Hetzner Online dedicated server plans are designed to provide a higher level of speed with an extremely stable network structure in our self-owned data centres in Germany. Coupled with best prices and outstanding support, we surpass customers' expectations around the world.

www.hetzner.info
info@hetzner.com

Kdenlive 0.7.3

LINUX FORMAT On the DVD

Video editing has become a de rigueur computing task, so **Andy Channelle** looks at whether Linux can keep up with iMovie and Windows Movie Maker...

In brief...

» A traditional non-linear video/audio editing suite. Also see *PiTiVi* and *Kino*.

Video is becoming increasingly widespread in our lives, but it seems as if Linux video editors have struggled to keep pace with those of other OSes. Sure, *Kino* is good, but it's not easy to get along with and *PiTiVi* isn't quite ready for its close up. Then, like a ray of sunlight bursting through dark cloud to a rousing score, *Kdenlive* appeared – a plucky young newcomer with aspirations to take on the likes of *Final Cut* and *Premiere*.

Kdenlive is a non-linear video editing application that takes the familiar timeline approach to editing. Adding and trimming clips works as you'd expect it to in other apps, which is also true for media management.

Initially, the screen is cluttered with different elements, but these can be opened, resized and closed depending on the kind of work you're doing. For instance, during capture we closed down the windows for the effects, transitions and timeline to concentrate on the capture process, then brought back the timeline and clip monitor during the rough cut phase.

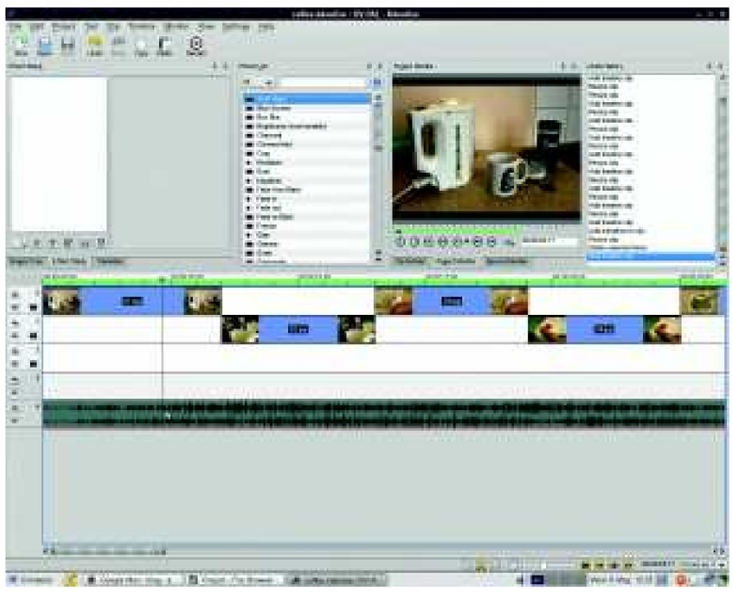

› ***Kdenlive*** **can manage as many audio and video tracks as your processor, memory and disks can handle.**

> **"Such flexibility makes it easy to concentrate on the task at hand."**

This flexibility, especially combined with full-screen mode (Ctrl+Shift+F), makes it easy to concentrate on the task at hand. Now you'd think all this configurability would mean *Kdenlive* is perfect for multi-monitor setups, but the software crashed repeatedly when running on this kind of system in our tests, even if it was only using one screen. Apart from this, however, the program is remarkably stable.

Keyframe power

Once clips have been imported, assembling your piece simply involves dragging and dropping them on to the timeline and using the Razor tool to cut or trim your media. It's an intuitive process and especially worthy of note is the Zoom slider, which provides a way to move from the project's overview to fine detail quickly. The right-click contextual menu is also well used, so you'll seldom need to use the menu bar.

Once you have a cut on the timeline, the software has a large number of effects that can be used to add a little style to the piece. Again, these can be dragged and dropped on to the timeline and then stacked or re-ordered using the Effects Stack window. Effects are divided usefully into Video, Audio and Custom, each with a set of tweakable options. There are also four transitions on offer. This may not sound like a lot, but the Luma transition takes care of a large number of dissolves, fades and wipes, while the Compositor handles split frames, on-screen graphics and so on. These effects can be used with keyframes, which makes it possible to create very sophisticated productions. Also, the live rendering of effects worked flawlessly, even when stacked.

In short, this is a brilliant video app. It's easy to get working and intuitive to use, but packed with enormous power. So, whether you're a budding film maker or just want to get your videos on to DVD, this is a great choice. LXF

Features at a glance

Easy interface
Editing audio and video is intuitive, but that's not at the expense of power.

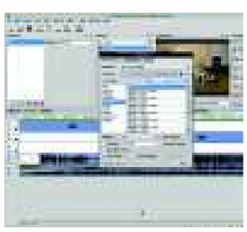

File formats
Output options cover almost every format you're likely to encounter.

LINUX FORMAT Verdict

Kdenlive 0.7.3

Developer: Kdenlive Team
Web: www.kdenlive.org
Price: Free under GPL

Features	9/10
Performance	9/10
Ease of use	7/10
Documentation	6/10

» *Powerful and accessible, which means it surpasses other platform's entry-level offerings with ease.*

Rating **8/10**

Ulteo OVD 1.0

LINUX FORMAT On the DVD

Gaël Duval's latest effort, Open Virtual Desktop, provides a virtual desktop in a browser. **Mayank Sharma** wonders if it can live up to its pedigree.

In brief...

» Delivers a virtual desktop running Linux and Windows apps. Proprietary alternatives include *VMware View* and *Sun VDI*.

Gaël Duval has been a busy bee. The man who brought us the user-friendly Mandrake (later Mandriva) was laid off in 2006 by the company he founded, but has since begun working on a project called Ulteo.

That brings us to the first usable Ulteo-branded offering, *Open Virtual Desktop* (*OVD*) – a thin-client system that uses one or two server machines (one for Linux, one for Windows apps) to serve a number of thin clients via a virtual desktop running inside a browser. And what's really cool is that this can be used on any host OS.

Although *OVD* doesn't do anything new (it's built on the concept of Virtual Desktop Infrastructure (VDI)), it has taken a hardcore enterprise back-end technology and brought it to the mainstream – much as Gaël once did with Mandrake Linux.

› **OVD delivers apps securely via SSH tunnels and can, according to the docs, run 20 concurrent users on 1GB of RAM from a multi-core box.**

Inside the duffle bag

The two main components that make up the Linux part of *OVD* are the application server and a Session Manager, both of which are open source and available for free. The Session Manager is a web-based console that administers the setup on the client machine, while the application server (as you might have guessed) serves up the Linux apps. Windows applications are instead served via a Windows agent that you must install on a separate copy of Windows Server 2003 configured with terminal services.

> **"OVD brings hardcore enterprise technology to the mainstream."**

This setup might sound daunting, but it isn't. There are binary packages available for the core Linux components and the installation process is well documented. You can install the server and Session Manager on the same machine and there's even a DVD that installs both on a customised copy of Ubuntu. The Windows bit is optional as well, so you can avoid it. Best of all, you can hook up *OVD* with a directory server – both Active Directory and LDAP are supported. Since all the processing is done on the application's servers, you can use *OVD* in a thin-client setup and a CIFS server for file storage.

Your UI

OVD comes with a set of pre-installed Linux apps, including *OpenOffice.org*, *Firefox*, *Adobe Reader* and *Gimp*. You can install more by specifying the name of the app in the Session Manager, which fetches it from a custom repository via *apt-get*. You can also associate individual apps to each user, or group similar ones together. The list of Windows apps, however, needs to be controlled via the terminal services settings from inside Windows Server.

Administering the setup is also done in the Session Manager, which enables you to configure a timeout period, display or hide desktop icons, and change keymaps and the UI language.

In our tests, the software easily served a couple of users with the server in a virtual machine running 128MB of virtualised RAM on a dual-core host.

For the future, Ulteo is promising features such as the ability to launch an app from a simple hyperlink without requiring a full desktop. For now, though, all that remains to be seen is whether *OVD* clicks with users. LXF

Features at a glance

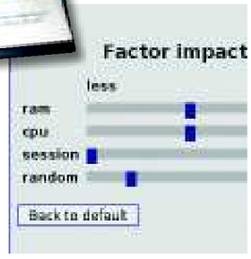

All bases covered
There's an in-built load balancer to select the best application server.

Collaboration
The controls help you share your desktop and collaborate with others.

LINUX FORMAT Verdict

Ulteo Open Virtual Desktop

Developer: Ulteo
Web: www.ulteo.com
Price: Free under GPL

Features	9/10
Performance	9/10
Ease of use	9/10
Documentation	8/10

» *Ulteo's OVD is a thin-client solution that's ideal for a corporate mixed-platform environment.*

Rating 10/10

Jets'n'Guns Gold

Faster, kill, faster, upgrade, kill, bomb, maim, injure. **Nick Veitch** relieves the stress of the supermarket with some old-fashioned entertainment.

In brief...

» Elaborately fashioned scrolling shoot 'em-up. See also *Project Starfighter*.

When you think of must-have games that should be ported to Linux, sideways-scrolling shoot 'em-ups probably don't feature heavily on anyone's list. The genre arguably reached its peak with *R-Type* in the murky, deafening arcades of the late 80s. However, it appears that in spite of being developed for a high-adrenaline coin-gobbling environment, that SSSEUs (to coin a rather ungainly acronym) still have some lure for the desktop gamer, and if you're going to rework such a beast, you might as well choose the best.

Jets'n'Guns combines all the classic game elements: a rousing soundtrack to put you in the killing mood, a simple control system, an element (OK, a whiff of a hint of one) of strategy, an array of weaponry not imagined since the dodgy dossier and some super-saturated smooth pixel graphic craftsmanship. It even has a great plot. No, actually, that isn't true, but it does have a story with some reasonable comic-style graphic panels between each level.

> **"There's plenty of tactical variation within the 42 levels."**

Shoot things, earn money, buy bigger gun is what it simmers down to. Actually, you can buy several guns, and this is one of the areas where *J'n'G* takes the whole idea a step further.

› **The shop is full of guns, and there's no cooling-off period.**

› **In the future, you need to fly modified biplanes and shoot things. A lot.**

Depending on your ship (you can upgrade that too) you can deploy multiple guns, missiles and bombs. You can have multiple weapons in storage, save them as sets and switch sets while you're playing – a nice twist, and useful where levels vary between ground targets, aerial squadrons and big boss-type targets.

I may close early today

There's plenty of tactical variation within the 42 levels of destruction, but most of them centre around blasting things, with a few puzzle asides as you try to pick the locks on cargo pods. There are secrets to be discovered, bonus levels to unlock and gory demises to be viewed, and all this is played out on some of the best artwork you've seen in a Linux shooting game (an OpenGL-capable graphics card is required). While the resolution is only 800x600, it's still pretty impressive.

The only real problem with this game, apart from it just getting a little too tedious when you have more than adequate firepower, is some of the more functional stuff. The mouse didn't seem that responsive at all, and trying to configure keys for a gamepad proved a little fruitless, as the direction pad wasn't accepted as an input (of course, joysticks are a whole different story on Linux anyway), and neither were a handful of other keys. Fedora users will most likely need to do some fiddling in the command line to get the game going, as it doesn't seem to find the correct OpenGL libraries on its own (it runs fine in Ubuntu on the same hardware though!)

With at least three hours of playing time, a real polished presentation and a liberal scattering of humour, this is certainly worth the money by commercial gaming standards. Yes, you have seen it all before, but probably not quite so well done. LXF

LINUX FORMAT Verdict

Jets'n'Guns Gold

Developer: Rake in Grass/ LGP
Web: www.linuxgamepublishing.com
Price: £15

Features	6/10
Performance	7/10
Ease of use	7/10
Documentation	5/10

» *Sideways-scrolling shoot 'em-up action to the extreme.*

Rating 7/10

EnergyXT 2.5

LINUX FORMAT On the DVD
» Trial version

Modular music composition is just another excuse for serial sonic sorcerer **Graham Morrison** to have several tracks on the go at once.

In brief...

» **Modular music creation with beat slicing, real-time effects and virtual synthesizers. Also consider *Ardour*, *Renoise 2.0* or *Rosegarden*.**

When we looked at *EnergyXT* in **LXF94** we were seriously impressed by the potential of this cross-platform modular music production environment. You see, rather than recreating the generic mixing console and arrangement views, *EnergyXT* enables you to tie together composition tools, sound generators, audio effects and recording modules into any configuration you choose.

Version 2.5 builds on this principle, but makes some concessions to usability. The package includes templates for MIDI sequencing, a multi-track recorder, a guitar project and the raw modular setup of the original, so you no longer have to get your hands dirty if you just want to create or record a tune. Click on the Drum And Bass template, for example, and a project with a drum machine and a bass synthesiser is loaded, defaulting to the note-editing view for both sound sources. After that, just click on the notes and press Play. You can create a complete track within seconds.

> **"Audio loops can be automatically sliced and stretched in real time."**

You can now completely avoid the manual wiring of the modular window if you wish. Effects, audio sources and even sequencers can be dragged from the resource panel into the main window, much as you can with *Ableton Live* on OS X and Windows. Another similarity is that audio loops can be automatically sliced and stretched in real time as you adjust the tempo. Samples, MIDI phrases and drum patterns can be dragged into the sequencer view, and even layered over each other to build increasingly complex patterns. Blocks of notes can be shifted left and right, and moved virtually to different modules quickly and efficiently.

› **Smooth scrolling, zooming, block copying and editing can be done from *EnergyXT*'s all-encompassing sequencer window**

Channel stripping

The Channel strip from the Mixer view can also be displayed alongside the tracks, enabling you to add in-line equalisation and drop-in effects, but things don't always work quite so smoothly. Dragging the Arpeggio module on to a synthesizer track, for instance, will add the module, but won't wire it in-between the notes from the sequencer and the synthesizer – you'll need to do this manually.

However, the biggest omission is the lack of LADSPA and DSSI support for native Linux effects and virtual instruments. Instead you're stuck with the effects bundled with the application, though it is possible to get a handful of Linux-compiled, free VST instruments working. Softening the blow are several new effects for version 2.5, including the retro-sounding Bit Crusher, a fantastic multi-mode filter, a compressor and a guitar amp model.

The internal synthesizer still sounds excellent, and can create all kinds of classic sounds, though it's a bit of a pain having to automate parameters like filter cutoff in the sequencer view: you have to first assign the cutoff to a MIDI control parameter and then edit this MIDI control in the sequencer view by double-clicking on a level for each quantised point within the sequence. It would be far easier to choose the parameter from a drop-down list and draw the automation with a Pencil tool. A new drum synthesizer and sampler splits sounds across a drum pad, and if you have the *Lame* libraries installed, you can now mix a project to an MP3. So despite the shortcomings in effects and instrument support, we can't think of another Linux audio app that's this flexible and as efficient to use. LXF

Features at a glance

Beat slicing
Loops can dynamically adjust to changes in tempo without altering the pitch of the notes.

Mixer view
The new mixing console lets you drag and drop effects into each channel.

LINUX FORMAT Verdict

EnergyXT 2.5

Developer: XT Software
Web: www.energy-xt.com
Price: €59

Features	8/10
Performance	9/10
Ease of use	7/10
Value	9/10

» *Fast, powerful, and lightweight. Would score higher if native Linux effects and synths were supported.*

Rating 8/10

heart internet

The UK's **Fastest Growing** Reseller Package

www.heartinternet.co.uk **Tel:** 0845 644 7750

Reseller Hosting

30 Day Money Back Guarantee

Only £29.99 Per Month

Reseller Features

- Host unlimited web sites.
- Linux & Windows accounts included. New
- Create your own web hosting packages.
- Disable, delete & edit accounts in real time.
- Fully brandable control panels and experience.
- Web hosting accounts instantly activated.
- HostPay shopping cart free.
- Generic webmail, SSL & nameservers or your own.

Ideal for Web Designers, Developers and Entrepreneurs

Start Your Own Web Hosting Company!

Web Hosting Features

- Unlimited web space & bandwidth.
- Unlimited mailboxes, forwarding & autoresponders.
- Unlimited databases, FTP & sub domains.
- Huge CGI library including WordPress & Joomla.
- Website builder & shopping carts.
- PHP, Perl, Ruby, SSL, Stats.
- Amazing control panel & file manager.
- Over 100 features for each account.

Too Much to List!
New Features Added All The Time

Check Out Our Website!

"Heart Internet supports their customers with market-leading hosting technology at great prices"

Suzi Perry
Technology TV Presenter

HostPay

FREE with our Reseller Package

Create Your Own Hosting Company

Amazing Features Including:

- Customer Management & Registration System
- Instant Activation For Your Customers
- Fully White Label Themes/Skins

What is HostPay?

HostPay is a highly customisable shopping cart & customer management system that is free with our Reseller Package.

- Supports PayPal, Google Checkout, WorldPay, Protx, Nochex & 2Co.
- Easy to set up and grows with your business.

For more information visit:

www.heartinternet.co.uk

Tel: 0845 644 7750 **Email:** sales@heartinternet.co.uk

MasterCard Maestro VISA Direct Debit

Prices exclude VAT

OpenOffice.org 3.1

Despite an uncertain future, the present is looking good for this prominent productivity suite. **Andy Channelle** takes a peek at its recent progress...

In brief...

» The foremost office suite for Linux. See also *KOffice*, *Abiword*/*Gnumeric* and Google Docs.

Undaunted by the myriad rumours of its death that have cropped up since Oracle bought Sun, the *OpenOffice.org* project has released version 3.1 of the productivity suite. This brings some much-needed smoothness to rendering, refinements to *OOo*'s interface and a small but significant boost to application load times.

Slow start times have been the project's weak point since its inception and that hasn't changed with this release. In our tests, *OOo Writer* – chosen because it is the part of the suite that we use most often – was ready to use from a cold boot in just over seven seconds on a 3GHz Pentium 4 machine with 1.5GB RAM. That sounds quick, but *Word XP* (running on the same machine via *Wine*) was two seconds faster. When relaunching, the times were four and two seconds respectively, so *OOo* is still being beaten by its rival. These are both kicked into a cocked hat, though, by *AbiWord*, which went from cold to usable in just 1.43 seconds. However, on a machine with a higher spec, the launch gap between *OOo* and *MS Office* was significantly smaller and both could be launched from a cold boot in under four seconds.

What's different?

Short load times might be good for bragging rights, but what's more important is how an application performs in use. Once loaded, *OOo* consumes just 19.5MB (ish) of memory, though this inevitably increases once documents are involved. In order to test heavy-duty word processing, we imported a large *Word* document containing a few high-resolution images and comments or amendments from four other authors. The application remained snappy throughout – there was no lag in selecting and editing text or moving images around the page, while panning and zooming was buttery smooth. Memory use jumped to 75MB, but the CPU load was fairly consistent, peaking at about 10% when copying or pasting a large volume of text.

> **"There was no lag in selecting and editing text, or moving images."**

On more modest hardware, memory and CPU usage may be an issue when working on massive documents, but we haven't had to endure frustrating waits for an object to render or appear yet, which is a definite improvement on previous versions.

Complementing the tweaks under the hood, there's another change that makes a difference across the entire suite – the addition of smooth, anti-aliased objects that can be added to the page. Where previous versions would display jagged edges, the vector images and drawings created within *OOo* are now displayed with smooth edges, no matter how large you make them. This is great when using Fontworks in *Writer* or when adding extrusions to objects in *Impress* to provide depth. Of course, this doesn't affect how bitmap images are displayed, but if you use *OOo* for presentations or light desktop design, it will make a real difference to the work you produce. It also has a big impact on the output of charts or graphs from *Calc*, which are now rendered brilliantly.

Welcome new content

Along with the overarching visual update, real object dragging has been added to make it easier to arrange documents – shapes and images are drawn as you drag them, rather than leaving you to work with a dotted outline. It's a small change, but it makes a difference if you're fanatical about precision, and it looks good too.

The version we downloaded for Ubuntu also had a good selection of

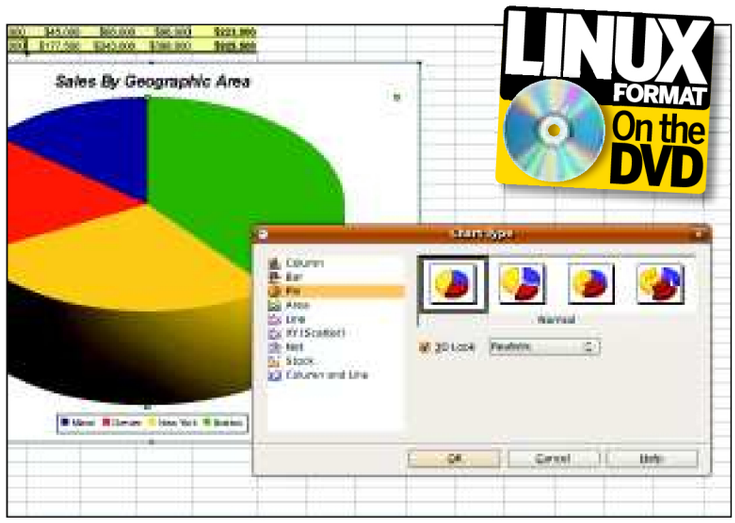

› **It's the improved, smoother representation of graphics such as charts and fonts that will have the biggest impact on casual users.**

Features at a glance

Smoother graphics
In-document anti-aliasing makes everything look a bit crisper – it's a great boon for presentations.

Chat back
The new comment system can turn a document review into a helpful conversation.

presentation templates to base your work on. They don't quite match the quality of the templates in Apple's *Keynote* or *PowerPoint 2007*'s new offerings, but they're a great place to start and, of course, there are plenty of user-generated templates on the internet if these don't suffice. *Impress* has also (finally) inherited the Zoom slider that was added to *Writer* in version 3, which makes it a doddle to navigate large, intricate slides.

Playing nicely

OpenOffice.org has always provided document properties that can be used to define elements of text embedded in a file, but version 3.1 has improved the process with Custom Properties. With these, it's now possible to add a range of metadata that makes document management a lot easier. For example, if you were drafting documents for a number of clients, your Custom Properties could contain information about the client, the person who commissioned you, deadlines and submission guidelines. The system contains a selection of predefined custom fields, but adding your own simply involves writing the property and then defining the content type (text, number, date and so on) and filling in the value. It's not going to affect every *OOo* user, but if you're creating a lot of documents or working in a collaborative environment, it could be quite handy.

Some of the other collaboration tools have been refined too. Version 3 of the software introduced margin comments into *Writer*, which made the process of multiple author editing similar to that in *Word*. This version adds comment replies that provide facilities for authors to comment on other author's comments, turning the editing process into a fluid, conversational affair. In our tests, the software also imported comments and tracked changes in *Word* competently, and exporting them was fine too.

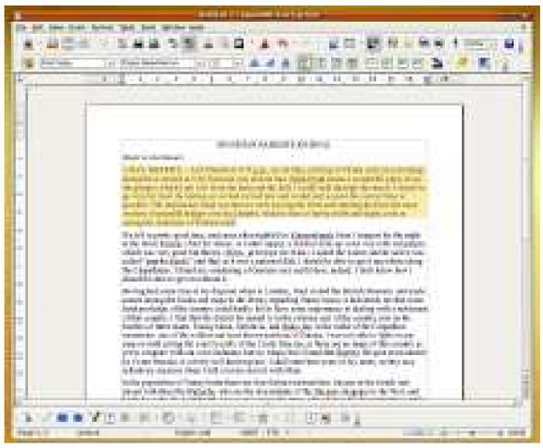

› Coloured highlighting improves usability a bit and looks far better than the old reverse-video effect.

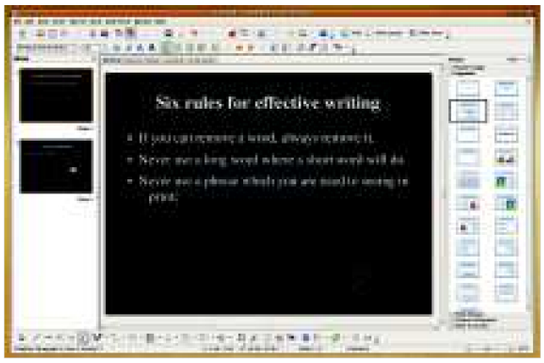

› A number of attractive templates are included in this version, and it's easy enough to build your own.

Graham says...

"I can't get excited by this release. I feel that for too long, *OOo*'s progress has been hampered by its de facto status of '*Office* killer'."

If you often work on collaborative documents, you'll benefit from better file locking. This means if one user is editing a document, the same document can't be opened and updated by another user. Previously, this worked well if every user was using the same operating system, but the new version makes file locking more reliable across different platforms.

With every new version of *OOo*, import and export options for common formats, such as DOC and XLS, improves. Following that trend, *OOo 3.1* has better support for the latest XML versions of the formats, but it's still not something we'd choose to rely on. In our tests, the majority of DOCX, XLSX, and PPTX files caused a crash or failed to display data. However, almost all older formats work with with no problems, so we'd recommend that collaborators stick to those if they intend to share their data.

The future?

With Sun's absorption into the mass of Oracle, it's been suggested that *OpenOffice.org* will be allowed to wither on the vine. While it's true that in the past the software has benefited from large numbers of Sun developers, this support also involved the same slow process management that often comes with corporate software development. In short, this may be an opportunity rather than a threat.

For years, Novell's Michael Meeks and others have complained that Sun's short leash on the project was causing it to fall behind and that useful features or patches could find themselves stuck in approval hell for so long that they were superseded before they hit the main software branch. With Sun's hand removed from the tiller, this could be the perfect opportunity for Red Hat, Novell, Canonical or IBM (all have a vested interest in the success of the suite), as well as the dedicated bedroom hacker, to step up and bring new enthusiasm and vigour to the *OpenOffice.org* fold.

There may be pain ahead, but *OpenOffice.org* is an important bulwark in the promotion of free software and so it's unlikely that any of the major players will allow it to die. And if development slows to a crawl as a new consensus emerges, at least we have some good quality software to use in the meantime.

Still #1

So, while *AbiWord* will give you a few extra seconds of spare time every day and *KOffice* might give you a better sense of community, *OpenOffice.org* should be your first choice of free software office suite. Where it really shines is in its compatibility with the still-dominant *MS Office* and the quality of its output. The speed improvements are welcome, while the refinements in graphics rendering, cross-platform collaboration and updates to the comments/changes tracking system will make it even easier for fans of *OOo* to work among themselves and with users of other suites.

The weakest part of the suite is the database, which isn't as stable as the rest of the software. It does perform a useful function in managing data sources for other parts of the package, but feels too complex for casual users.

Finally, the parity of features that *OOo* offers across the three main operating systems means that it's ideal for those migrating from Windows to Linux, keeping the software familiar while the OS changes. All in all, *OOo 3.1* is highly recommended. LXF

LINUX FORMAT Verdict

OpenOffice.org 3.1

Developer: OpenOffice.org project
Web: www.openoffice.org
Price: GPL

Features	9/10
Performance	8/10
Ease of use	8/10
Documentation	7/10

» *A marvel of free software development and an essential download for all Linux users.*

Rating 8/10

Learning Rails

This book should stop **Andy Hudson** going off the Rails.

Learning Ruby has proved to be a popular book since its release and now O'Reilly has delivered a text that perfectly complements it. *Learning Rails* aims to take you from no knowledge of Rails to being able to build applications using the Rails framework. Before you go near this book you'll need to make sure your HTML knowledge is up to scratch, and you'll need to have a programming mindset – this is essential if you're going to follow the many code snippets that appear throughout.

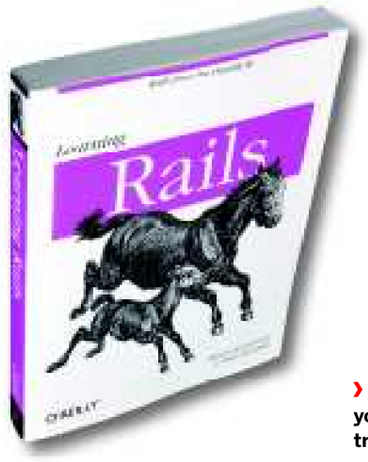

› Read this and you'll be on the right track with Rails.

Unusually for a programming book, *Learning Rails* focuses on the graphical results of your code first and foremost, teaching you to appreciate what your users are going to see over how you code it; this approach serves the book well, as it focuses on the end results. The authors try to strike a fair balance between telling you the basics and giving you enough detail for you to want to learn more, and the book never tries to move too quickly, with each chapter being rounded off with a short quiz to hopefully embed what you've learned.

Of particular note is the penultimate chapter, which deals with securing, managing and deploying your Rails applications. The authors give you a heads-up on some of the more common types of attack, along with details of how to prevent them from being exploited.

Learning Rails is completed with a series of brief introductions to Ruby, relational databases and the all-important regular expressions. Each of the appendices are intended to only give you a flavour of each topic, rather than being something to replace a dedicated source of information, but the level of detail is enough to enable you to understand the basics. LXF

LINUX FORMAT Verdict

Learning Rails

Author: Simon St Laurent & Edd Dumbill
Publisher: O'Reilly
ISBN: 978-0-596-51877-6
Price: £24.99
Pages: 398

» *A valuable introduction to Rails and the ideal companion volume to* Learning Ruby.

Rating 8/10

Blender for Dummies

Graham Morrison initially thought that this title was an insult...

We try to maintain a certain air of technical aloofness at LXF Towers. When people look over our collective shoulder, we'll open the command line and type **cat /dev/urandom**. It never fails. Which is why there was shock and disbelief when someone noticed the distinctive yellow of a Dummies book lying on the desk. "No, it's not a gift from an angry reader," we quickly said, "It's for review." Despite its power and capabilities, *Blender*'s bewildering GUI is a tough challenge for anyone, us included, and a Dummies guide to getting the most from it seems like a great idea.

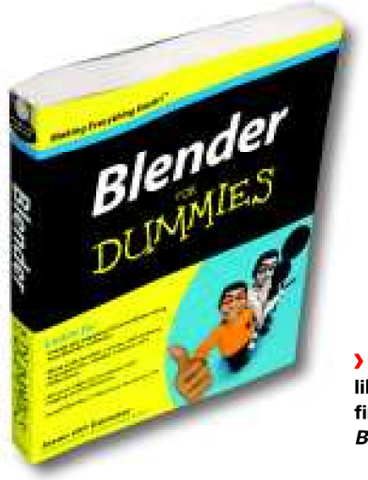

› Everyone feels like a dummy the first time they give *Blender* a try.

Which *Blender for Dummies* mostly is, in execution. This book is full of the two-sentence descriptions and the cropped screenshots that have made the other Dummies titles so popular. It's not dumbed down, and it's not trying to be the magic pill that suddenly makes *Blender* easy to use either. There's still plenty to read.

The book spends most of its time trying to get the reader into the *Blender* mindset, in the hope that everything else will fall into place. It also tries hard to explain, or excuse, some of the design decisions behind the *Blender* interface. You have to wade through a couple of heavy chapters on keyboard commands and the various *Blender* panels before you get to create anything interesting. But the later chapters on meshes, textures, sculpting and animating are particularly helpful. There are no step-by-step tutorials, but the words do carry plenty of authority and anecdotal examples. The book constantly reminds the reader that they're not alone in finding the world of *Blender* confusing, which by the end of the book is its greatest success. LXF

LINUX FORMAT Verdict

Blender for Dummies

Author: Jason van Gumster
Publisher: Wiley
ISBN: 978-0470400180
Price: £22.99
Pages: 408

» *An excellent, readable reference resource for those who are new to the world of 3D animation.*

Rating 9/10

Automating Linux and Unix System Administration

Andy Hudson likes nothing more than a quiet life; will this book help him?

If there's one thing any system administrator prays for, it's an easy life. Apart from having to deal with annoying developers, ignorant users and management red tape, it's a fact of life that you'll never be able to keep on top of everything. This book aims to give you some ointment for your pain by teaching you how to carry out common and repetitive tasks using the *cfengine* suite. *Automating Linux and Unix System Administration* is pegged at the far end of the Apress roadmap, meaning that only experienced sysadmins should read this; it's certainly overkill if you're only looking after one or two machines, for example.

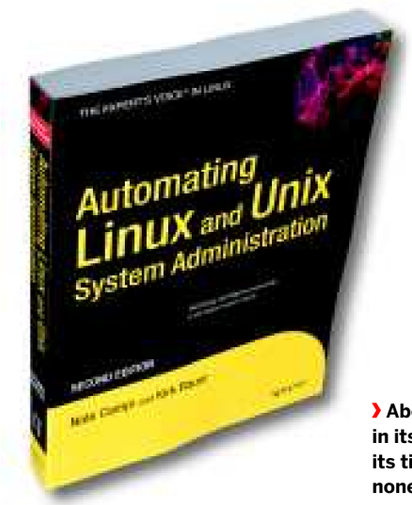

› About as succinct in its approach as its title, but helpful nonetheless.

The early chapters deal with some basic fundamentals, but they're touched on only briefly before it moves on to the practicalities of automation. Followed by a brief introduction to configuring SSH to work as needed, you're given an overview of *cfengine* as well as the configuration of the *cfagent.conf*. From there on in, you're shown how to create an infrastructure that will support *cfengine* along with a useful chapter on automating installs of Debian, Solaris and RHEL. Once you're past this, you'll start to work towards building an automation solution that will genuinely help you. The book finishes with a couple of appendices, the first of which is designed to give you primers in the basic toolkit that any sysadmin should be familiar with.

It'll take you some time to get into this book, but the benefits are clear; one annoyance we found was the insistence of the authors to print out entire configuration files, which would have been better placed on a companion website, but this is a minor gripe. LXF

LINUX FORMAT Verdict

Automating Linux and Unix System Administration

Author: Nate Campi and Kirk Bauer
Publisher: Apress
ISBN: 978-1-4302-1059-7
Price: £39.49
Pages: 400

» *A useful guide to this complicated area, but quite long-winded and with a propensity to print every conf file.*

Rating 7/10

Distrowatch

This month's roundup of news from the vaults of Linux distro development.

SUSAN LINTON is owner/ operator of tuxmachines.org.

Tainted ideal

The first week in May I saw something I never dreamed I'd see in the open source community. Clement Lefebvre, lead developer of Linux Mint, posted on his distro's blog to request that anyone from Israel or anyone who supports them in the Palestine/Israel conflict refrains from contributing to or using Linux Mint.

It's a great shame, because the open source and Linux community was the one place where race, creed, religion, gender, political party, country of origin, sexual orientation, and, for the most part, monetary worth was truly irrelevant. No matter what went on in the world around us, we remained friends and coworkers with a common interest. The open source community had achieved by accident what activists and governments have been trying to create for decades – a level playing field. Well, until this post from Clement.

It doesn't matter what my views are on that age-old conflict, because I'll still feel resentful when I use what had been one of my favourite Linux distros. Not because of Clement's political beliefs, but because he dragged those beliefs into our world. The only political debates appropriate in our community deal with digital rights management, patents, open source adoption and net neutrality.

lxf.distrowatch@futurenet.com

Spring clean

Mandriva 2009.1 The Spring version of this beginner-friendly distro brings some seasonal excitement.

One of the big players in desktop Linux released its latest and greatest recently amidst a deluge of Ubuntu news. Unfortunately, that meant Mandriva didn't receive the same amount of coverage, but early reviews all sang its praises. One blogger even went as far as saying Mandriva , "kicks Vista/ Windows 7 back to **/dev/null**".

The release builds on some of the improvements in 2009.0, such as the drivers for Atom-based netbooks and faster boot times. Mandriva 2009.1 also features KDE 4.2.2, Gnome 2.26, *Xfce 4.6*, *X.org 7.4*, *X.org Server 1.6*, Linux kernel 2.6.29 and *GCC 4.3.2*. Meanwhile, the list of applications includes *Firefox 3.0.8, OpenOffice 3.0.1* and *Gimp 2.6.6*.

The Mandriva team enhanced their collection of configuration tools for this release as well. The *Mandriva Control Centre* has received a GUI revamp and the MSEC security package has been heavily modified and redesigned. Among the many improvements are support for plugins, custom security levels, desktop notifications and advanced logging facilities.

Mandriva 2009.1 has received refined networking capabilities too. These include new firewall options, enhanced network monitoring, better 3G support, a reduced memory footprint and improved wireless support. In addition, Suspend, Hibernate, and Resume are also much more responsive in this release.

› ***KWin* effects, such as windows that explode when closed, can be activated in the All Effects tab of the new KDE Control Panel.**

> **"Mandriva 2009.1 is among the first distributions to offer a KDE 4.2.2 desktop."**

But it's not just the desktop that's improved – the installer has received attention as well. Some of its new elements include a Hardware Detection tool that provides some low-level hardware information right from the boot menu. Another extremely handy feature enables users to view the contents of their partitions before making any disk changes. Ext4 is offered as a filesystem and this release can be installed directly on an ext4 partition. Finally, the partitioner can now handle partitions that start after the first terabyte of your hard disk.

Mandriva 2009.1 is among the first distros to offer a KDE 4.2.2 desktop. The default folder view removes the folder widget, so icons can now appear directly on the desktop, and this familiar arrangement eases the transition from KDE 3 to KDE 4. Finally, the *KWin* desktop effects can even rival the dominant (and enabled by default) *Compiz Fusion*.

Mandriva is available in a 4.4GB free download containing only open source software, 700MB Gnome and KDE live CDs that contain some proprietary code, and the commercially available Powerpack.

Mandriva 2009.1 isn't just a bugfix release. It offers plenty of goodies in a stunning, polished and stable framework.

www.mandriva.com

Jiggy with Jaunty

Ubuntu Canonical's latest has arrived, but can it live up to the hype?

The latest version of Ubuntu has received a storm of attention, ranging from gushing exaltations to damning condemnations. Some of the more negative comments arose from 9.x teething issues, such as upgrade failures and graphics-related problems.

It seems that no graphics manufacturer was spared: reports of black screens, locked desktops and poor performance abounded. A new kernel update seemed to help with some Intel-based performance issues, while those experiencing Nvidia and AMD troubles turned to help from volunteers at the Ubuntu forums.

Aside from this, though, the release brought many new features, although some users were disappointed that the distro's new look has yet to materialise.

The list of goodies starts with vastly improved boot speeds, including enhanced performance when using Hibernate, Suspend and Resume. The ext4 filesystem is also available, along with better 3G and wireless support. In addition, intelligent switching is among the feature set, and private and hybrid cloud computing environments can be created with the server edition of *Eucalyptus*. Some smaller features include the new Notifications applet, added themes and improved multiple monitor handling.

But Jaunty isn't revolutionary –it's unlikely to lose fans or woo Ubuntu cynics.

www.ubuntu.com

› Tweaks to the installer make Ubuntu friendlier for newbies and pros alike.

In tiny packages

SliTaz GNU/Linux This small distro gives a full desktop experience.

One of the coolest little distributions available today is the diminutive SliTaz GNU/Linux. Weighing in at just 30MB, it manages to provide a great desktop experience, along with good applications and utilities. Version 2.0's recent launch saw a rather quiet reception, but the lack of much buzz for this fledgling distribution isn't an indication of its worth.

SliTaz comes as an installable live CD that takes advantage of the small and fast *Openbox* window manager. It includes software such as *Firefox 3.0.8*, *AlsaPlayer*, *MTPaint*, *GpicView*, *Ghost In The Mail*, *Grsync*, and *Transmission*. In addition, SliTaz maintains a repository of 1,400 software titles that are easily installed with its *Tazpkg* package manager, which can be invoked from the command line, or through a tidy graphical front-end.

Under the bonnet is Linux kernel 2.6.25, *X.org 7.3*, wireless drivers and tools. SliTaz also comes with several config tools to customise your system and desktop settings. Useful monitors and applets populate the panel and, helpfully, the documentation is on the desktop.

SliTaz can be used on a wide range of systems, because it only requires a i486 processor and 256MB RAM. It can be installed on a USB memory stick and you can customise your own ISO. Overall, it's an impressive distro in a tiny package.

www.slitaz.org

› The easy-to-use *Tazpkg* front-end works very much like the ever-popular *Synaptic*.

Future hardware

Netbooks were rumoured to be losing some of their shine last autumn, but along with the economic crash came renewed interest in the tiny form factor. The only downside is that Linux now seems to be available on fewer machines while Windows XP is being used more. Will that affect Linux users' choices, and what type of computer will characterise their next purchase? We asked **tuxmachines.org** visitors this question and here's an overview of what 597 of them said:

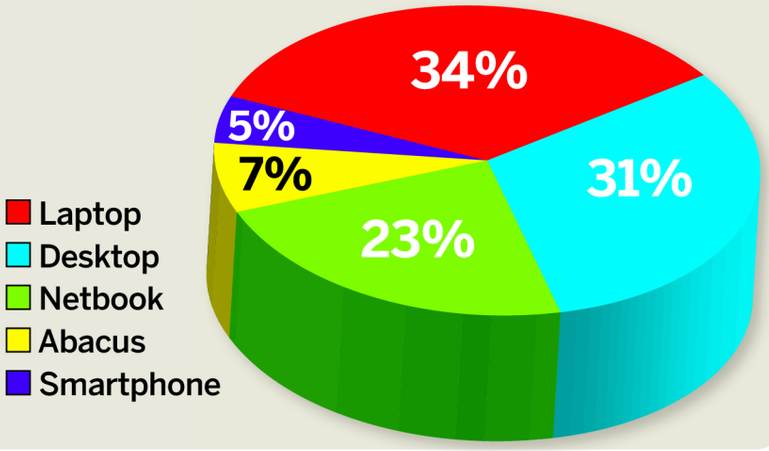

Hit list

The ten most visited distro pages on DistroWatch.com, 29 April–28 May 2009 (average hits per day)

	Distro	Number of hits	
1	Ubuntu	2,057	‹›
2	Fedora	1,456	↑
3	Mint	1,432	‹›
4	OpenSUSE	1,256	↓
5	Mandriva	1,217	↑
6	Debian	957	↓
7	Puppy	823	↑
8	Sabayon	809	↑
9	CentOS	585	↑
9	Kubuntu	585	↑

» **DistroWatch.com** monitors the popularity of distributions based on the number of visits to each of its distro-specific pages. While these figures don't represent the actual install base, it's an indicator of which distros were hot during each specific time period. LXF

Pick the perfect netbook

We locked **Mike Saunders** in a room with eight netbooks, a week's supply of oxygen and one mission: to find the ideal machine for every type of user.

We're not rabid Microsoft-bashers here at *Linux Format*, but we always have a chuckle recalling Bill Gates's tablet PC prediction at Comdex in 2001. "Within five years I predict it will be the most popular form of PC sold in America", said the world's most famous IT billionaire – yet he wasn't even close. Tablets are still largely regarded as novelties and confined to a few niche market segments.

What has taken everyone by surprise though is the booming netbook market. When Asus released the Eee PC 701 most pundits thought it a pointless exercise, but its skimpy hardware proved more than adequate for the tasks that most people do on a regular basis. Light web browsing, office work, solitaire on the train – the Eee did a good job, especially at its £200 price point, and was soon followed up by models with larger screens and keyboards to mitigate some initial gripes.

Major players such as Dell, Lenovo and Toshiba are now in the game, and while these machines mostly fit into the low-cost (sub-£300) category, choosing the right one is essential while we're all watching our pennies. They're all capable of internet browsing with Wi-Fi and office work, but they can vary drastically in key areas as we'll see.

Our test criteria

We've brought together all the netbooks we could get hold of for a comprehensive test. We're looking at:

» **Performance** All but one of the netbooks are based on the Intel Atom 1.6GHz CPU and 945GME graphics chip. But other components come into play, especially the storage and the wireless reception strength, so we're putting particular focus on these aspects.

» **Usability** The most important aspect of a netbook. It doesn't matter if it looks wonderful if the keyboard is far too cramped, or the trackpad is rubbish.

» **Build quality** You shouldn't need to baby your netbook. You want to chuck it in your bag, use it everywhere and not worry about it taking a bump or two.

In order to make our benchmarks fair, and because we know that most regular Linux users prefer to install their own distro, we'll install Ubuntu 9.04 Netbook Remix on each machine that supports it. Let's get started then...

Contents

Reviews:

Acer Aspire One p39
ASUS Eee 1000 ..... p40
Dell Mini 9 ..... p41
Elonex ONEt p42
Lenovo IdeaPad p43
LG X110 ..... p41
MSI Wind U100 ..... p44
Toshiba NB100 ..... p45

Acer Aspire One A110

Price: £139 (www.simplyacer.com)
Web: www.acer.co.uk

Recommended for: Budget netbook shoppers, internet-only use.
Not recommended for: Gaming, demanding tasks.

Outside of the Eee series, the Aspire One models from Acer are the best-known netbooks on the market. Currently two models of the 110 range are available, one featuring flash (SSD) storage and the other a more traditional hard drive. Here we're looking at the 8GB SSD model, which has recently dropped to an extremely tempting price point – if you look around online you'll find it for a recession-mocking £139.

Cosmetically Acer has made a big deal with curves on the Aspire One. The lid and underside bend neatly to a point, and on opening the machine you're greeted with a bulbous section beneath the screen, housing status lights and the battery. There's a very visible gap between the hinges that hold the screen on to the body of the machine, giving an initial impression of weakness. However, the Aspire One is generally very sturdy with only a little flex in the screen.

And that screen is glossy all the way. Whether this is a big deal to you is a matter of personal choice, but some users simply hate glossy screens. Position yourself incorrectly and you'll have all manner of light sources glaring back at you, obscuring the stuff on your screen, which isn't too difficult to fix in an office/home environment – but outside, you can't get rid of the big flaming orange thing in the sky. Conversely, colours are strong and bold, and it's an excellent display given the price of the machine.

› **The model with a hard drive is a tad thicker and a lot more expensive than the SSD version on test here.**

Port plethora

The left-hand side holds an SD card slot, USB port, Ethernet port and VGA out, while the right side has an extra pair of USB ports, another SD card slot, the Kensington security socket and headphone/microphone ports. On top, the keyboard is great – good-size keys, a chunky Enter button and great overall feel. The cursor keys are a bit small, but at least there are dedicated Page Up/Down buttons. We're not so chuffed with the trackpad though: it's very small, and to salvage space Acer's designers have positioned the click buttons on either side of the trackpad. Now, you can retrain your muscle memory to deal with this, but we still prefer the normal layout.

On the software side, the Aspire One is bundled with Linpus Linux Lite, a Fedora-based distro that's supplied with the apps you'd expect – *Firefox*, *OpenOffice.org* etc. It's a decent distro with a great boot time (22 seconds), but quite locked-down and dated, so many users opt for a 'full' distro instead. Ubuntu 9.04 Netbook Remix runs very well, with Wi-Fi, webcam and sound working out of the box.

"Cosmetically Acer has made a big deal with curves."

There's one let-down though, and it's the SSD performance. It's cripplingly slow at times, especially when writing data. Linux (like any modern OS) builds up file write operations to save in one big batch, and on the Aspire One this manifests itself as annoying intermittent lock-ups. Just switching categories in the Ubuntu Netbook interface took 11 seconds at one point while we were using it because the SSD had some activity to finish.

Ultimately, this isn't a problem if most of your work is online: it's a great little web-browsing machine and handles YouTube with ease. The fan kicks in quite a lot and the speakers are tinny, but that's not a concern at this price. If you plan to do work on the go, though, or play games, you'll find the stuttering SD very painful – and with the hard drive model's price tag around the £220 mark, it's not so much of a bargain. In that case you'll want to consider the others on test here, such as the Eee PC.

› **Linpus Linux, the default OS, keeps a lot of its functionality hidden from the user.**

LINUX FORMAT Verdict

Performance	5/10
Usability	7/10
Build quality	8/10

» *Outstanding value and great for mobile internet use, although the slow SSD is horribly frustrating.*

Rating **7/10**

Asus Eee PC 1000

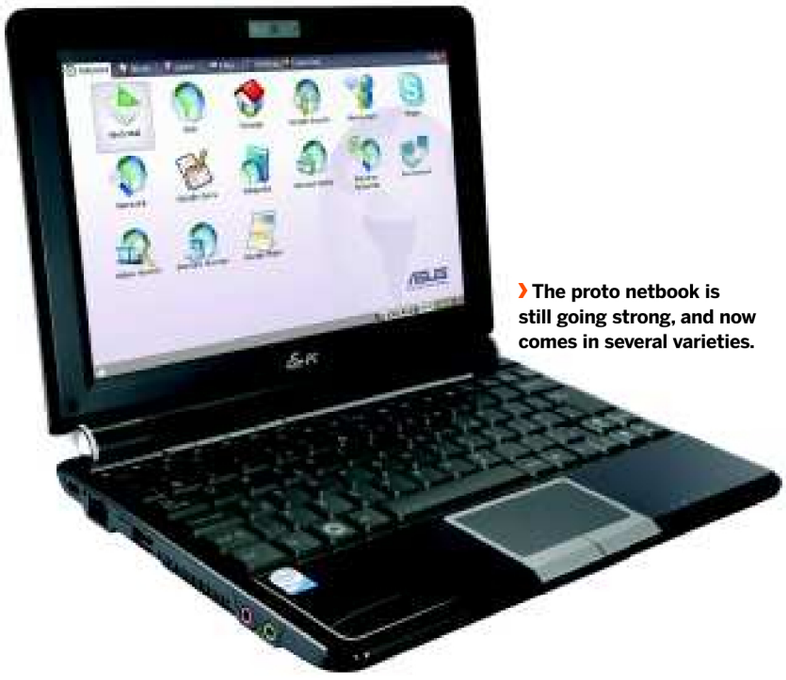

› The proto netbook is still going strong, and now comes in several varieties.

Price: £249 (Amazon)
Web: www.asus.co.uk

Recommended for: People looking for usability close to that of a full-size laptop
Not recommended for: Anyone looking for all-out portability

Here's the big one, in two senses of the word. Not only is Asus the best-known name in the netbook world, with a bewildering range of models on the market, this particular Eee 1000 is also the biggest machine in this test. It's almost approaching a regular laptop in size terms, and it's pretty heavy at just under 1.3kg too. So that's the most important thing to consider: if maximum portability is what you're seeking in a netbook, the Eee 1000 isn't for you – have a look at our reviews of the Toshiba and the Dell Mini.

> **"The colours are vivid and the image razor sharp."**

With the Eee's chunkiness, though, comes a very convincing build quality. Plastics don't creak when pressed hard, the screen is very rigid and almost impossible to flex, and the keyboard doesn't rattle around either. Asus's original 701 model was more than capable of taking a few knocks, so we're glad to see that the company has maintained the focus on a solid design. Part of the weight comes from the whopping standard six-cell battery that also raises the thickness of the unit – Asus claims up to seven hours battery life, and we'd say that's not far off as long as you're not pummelling the CPU.

Prepare to celebrate if you hate glossy screens: the 1000's is matte. Nonetheless, the colours are vivid, the image razor sharp and it's one of the best screens in this test. It has a 1.3MP webcam, while on the top-side of the keyboard there are buttons for switching video modes (ie to a separate display), launching *Skype* and changing the performance level. We're mostly thumbs-up with the 1000's keyboard – a great size, huge Enter key and full-size cursor keys. The only let-down is the small, rattly right-hand Shift key.

Given the dimensions of the machine, Asus has been able to plop on a satisfyingly large trackpad with big buttons to top it off. Its behaves rather strangely in the default Xandros Linux distro; it's quite floaty, as if the pointer moves a few pixels further when you release your finger. This is something that our friends on sister mag *PC Format* have seen with their own Eee, too, so we know it's not limited to our review model.

The usual suspects

On the left is the Kensington lock, Ethernet jack, air intake, USB port and headphone/microphone jacks, while the right-hand side is home to the power port, SD/MMC card slot, VGA out and two more USB ports. As mentioned, the Eee is supplied with Xandros Linux sporting bold and chunky desktop icons. It's a pretty friendly Linux flavour for newcomers, but the stock software selection is woefully dated (*Firefox 2*!?) and the update manager entertainingly crashed when we tried to check if there was anything new.

As expected, Ubuntu ran well, though it didn't detect the built-in Bluetooth module. The 40GB of SSD drive storage in our machine was split across two drives, one of 8GB and the other of 32GB – we installed Ubuntu on the 8GB drive for our benchmarking. Given the Eee's size, we'd expect there to be plenty of room inside for good CPU ventilation, but annoyingly the fan kicked in a lot in our tests, and has quite a loud sound too. A bit of a shame, as otherwise it's a great experience.

LINUX FORMAT Verdict

Performance	8/10
Usability	9/10
Build quality	9/10

» *Big and heavy, yet great value, superb battery life and the closest thing to the usability of a laptop.*

Rating 9/10

Other netbook distros

While Ubuntu Netbook Remix is the most prominent netbook-focused distro, there are plenty of others. Mandriva has made a push towards wider netbook compatibility in recent releases, while Debian has a stash of info at **http://wiki.debian.org/DebianEeePC**. Some others to look out for:

Foresight Linux Mobile shares the same front-end as UNR.

» **Easy Peasy** Ubuntu spin-off that includes proprietary extras such as Flash and *Skype*. **www.geteasypeasy.com**.

» **CrunchEee** Another Ubuntu derivative, geared towards more experienced users **www.crunchbanglinux.org**.

» **Foresight Linux Mobile** Uses the unique *Conary* package management system **www.foresightlinux.org/mobile.html**.

» **Pupeee** Lightweight distro, well suited to lower-spec devices such as the Asus Eee 701. **www.puppylinux.org/wikka/Pupeee**.

Dell Mini 9

Price: £228
Web: www.dell.co.uk

Recommended for: Portability and silence seekers
Not recommended for: Unretrainable Fx key users

The Mini 9 is nicely portable, though this takes its toll on the size of the keyboard.

Can a netbook really be silent? Well, the Mini 9 doesn't include a hard drive and it's also entirely passively cooled. This is quite a feat given that it has the same CPU as the Eee 1000, and that laptop is a loud beast at times. Still, the downside to this is that the machine gets quite hot underneath when the CPU is stressed.

The Mini 9 is not adorned with extra fancy metal rims like the Eee or Aspire One. The top of the lid is a haven for fingerprints, while the screen is highly glossy and reflective, which may give you some grief if you're working outdoors. On the left-hand side you'll find the Kensington lock, power plug, two USB ports and the SD/MMC card slot; the right side is home to the headphone/microphone jack, another USB port, a surprisingly small air vent, VGA out and Ethernet.

Dell has made some brave decisions here, most notably removing the Fx keys from their usual habitat of the top of the keyboard. Instead you have to press Fn+A to the semicolon key on the keyboard, which requires some mental retraining.

Another oddity is the right-hand Shift key, which is two keys in from the left, past the back tick (`) key and up cursor key. But these things aside, the keyboard is quiet and robust, with keys of the same width as those on the Toshiba NB100, but with extra height that makes it considerably more usable. The trackpad is good too, reaching right up to the keyboard.

Dell ships the Mini 9 with Ubuntu 8.04 LTS, and Netbook Remix 9.04 ran like a charm on the machine, with the snappy SSD giving it the second-fastest boot time in our benchmarks.

Battery life is good too, though the Wi-Fi reception was the second worst of the bunch, reporting only a 40% signal strength where others would get 70% upwards. Still, for £228 it's good value, and you can knock that down to £199 if you omit the webcam.

LINUX FORMAT Verdict

Performance	8/10
Usability	6/10
Build quality	9/10

» *A well-designed, silent machine with a fast SSD, but with poor Wi-Fi reception and a cramped keyboard.*

Rating 8/10

LG X110

Price: £284 (www.handtec.co.uk) **Windows only!**
Web: http://uk.lge.com

Recommended for: Power users, Windows dual-booters
Not recommended for: Anyone who doesn't want to give Microsoft money

We'd upgrade the battery, but otherwise the LG X110 performs well.

Where many of the machines in this test opt for mean-looking black and blue shades for their keyboards and cases, this LG mixes up soft shades of white and silver. The lid is very understated: it's almost entirely white, save for a small LG logo on the cover.

On the hardware side, it's very clear that the X110 is a rebadged MSI Wind (reviewed later), with an almost identical port layout and the same huge air vent on the left-hand side. But there has been more effort on the appearance: whereas the Wind's headphone and mic jacks are marked with the usual green and red rings, on the X110 they retain the colouring of the case. Additionally, the X110's VGA port is black rather than blue, further enhancing the professional looks.

The biggest differences between the X110 and the Wind are in the keyboard and build quality. The X110's is superb, quiet and with all the keys in their expected places – note that the cursor keys are not full-sized though. And while the X110's hinge is a bit weaker than, for instance, the Eee or IdeaPad, the rest of the unit feels very tough, with a better choice of materials than the Wind. Our review unit came with a three-cell battery that performed poorly in the tests (page 46) – we recommend splashing out an extra £35 on a six-cell version.

Despite the similarities in hardware, the X110 has a different wireless chip and no Bluetooth. We're not going to dwell on this machine for too long though, simply because it's only available with Windows. If you regularly use Microsoft's OS and want to have a dual-booting netbook, it's brilliant – very well designed and a good performer. If you've got no use for Windows but like the other aspects of the machine, consider the mostly-similar Wind.

LINUX FORMAT Verdict

Performance	8/10
Usability	9/10
Build quality	9/10

» *An attractive, powerful, quiet and robust machine if you need to dual-boot with Microsoft Windows.*

Rating 8/10

Elonex ONEt

Price: £119 (Elonex)
Web: www.elonex.com

Recommended for: Kids, hackers who hate x86
Not recommended for: Anyone else

› **The ONEt squeezes an impressive amount of juice from its two-cell battery.**

Even if you've never heard of the Elonex ONEt before, you've probably seen it in a different incarnation. You see, in China there's a gigantic factory pumping out netbooks with mildly varying designs but identical specs: 400MHz CPU, 128MB RAM and 1 or 2GB of flash storage. You'll see it in Maplin as the Minibook, or elsewhere as the Skytone Alpha 400.

Elonex offers the ONEt in a trio of colours – black, green and pink – and right from the first time you pick it up, you know that it's tough. The firm frame and thick plastic make the ONEt feel satisfyingly robust. It's the smallest netbook on test here, just a smidgen narrower than the Toshiba, and this becomes a problem with the keyboard (especially the tiny cursor keys) and the nanoscopic trackpad. To conserve space, the latter has buttons on each side *à la* the Acer Aspire One.

> **"From the first time you pick it up, you know it's tough."**

On the right-hand side is a pair of USB ports, while the left-hand side holds an SD card slot, headphone jack and microphone input port. Round the back is another USB port, an Ethernet port and the power socket – in all, not a bad range of pluggable holes given the low price.

So far, so good-ish: it's tough and well equipped. What about the software? Well, the ONEt is unique in this roundup in that it doesn't have an x86-compatible CPU. While all other netbooks are theoretically capable of running any PC OS from Windows 3.1 to Fedora 11, the ONEt and its brethren use a CPU based on the XBurst architecture. XBurst is a variant of MIPS, the CPU family famously used by SGI in its beefy-looking Indy and Octane workstations. Generic Linux distros and PC operating systems are a no-go on the ONEt then, which may dissuade some buyers. But far more problematic is the processing power. The XBurst chip simply lacks the muscle for non-trivial internet tasks, stuttering painfully on websites like BBC News and freezing up completely with heavy JavaScript (such as Google Mail).

Limp software

This isn't helped by the software. The ONEt is supplied with a bespoke CELinux-based distro that includes *Firefox 2* and therefore its sluggish JavaScript engine, and official Flash isn't an option here, due to the MIPS core. The tabbed interface provides access to *AbiWord*, *Gnumeric*, four games (including *Pengupop*) and a locked-down file manager that makes sure you don't touch anything outside your home directory.

There's also a media player that helpfully flashes up the max recommended specs when you start it (350x288 resolution at 128kbps). Power management is almost entirely absent – there's just a "bettery" meter, and the screen backlight doesn't switch off when you close the lid.

There are a few add-ons at **www.littlelinuxlaptop.com**, along with a community-built distro called 3MX, which is pretty cool in a geeky kind of way, but you'll still only want to run *Dillo* for web browsing. Still, 3MX is proper Linux and gives you a sense of control back.

Overall, the ONEt is sorely insufficient in the "mobile lifestyle" category; it's really awkward to use for web browsing. If Elonex dropped the price to £69 and crammed it with education software such as *GCompris* and *KGeography*, it'd be perfect for kids. But unless you've got a hankering to play around with a portable MIPS box, it's simply far too underpowered and not good value

LINUX FORMAT Verdict

Performance	2/10
Usability	3/10
Build quality	9/10

» *It might keep the kids busy for a bit but otherwise the £20 extra for an Aspire One is completely worth it.*

Rating 3/10

MS's market share

At the start of April, Microsoft proudly announced on its Windows Blog that it had snapped up 96% of the netbook operating system market (**http://tinyurl.com/msnetbookshare**). As Linux fans we all find that a little bit worrying – Canonical posted a good response (**http://blog.canonical.com/?p=151**) but there's no doubt that Linux is being sidelined in some parts of the market, as seen by the Windows-only netbooks in this test.

One intrepid member of Team LXF was in PC World a couple of months ago, and overheard the sales lady telling a netbook-seeking couple that "you don't want this because it's Linux and it won't run your software". Now, perhaps the couple had explicitly stated that they wanted a Windows netbook from the start, and her guidance was spot-on. But we wonder if shops and sales staff around the world aren't even giving Linux a proper chance – and, of course, if you sell a Windows netbook then you can also sell "security suites" and all sorts of other crack-over-papering add-ons...

Lenovo IdeaPad S10e

Price: £289 (www.microdirect.co.uk, three-cell battery version)
Web: www.lenovo.com/uk/en

Recommended for: Pro users who need an ExpressCard slot
Not recommended for: People working in quiet environments, thanks to the ultra-clicky trackpad buttons!

Lenovo isn't a household name when it comes to computers, but in business circles the company is well respected for its robust line of ThinkPad laptops (previously produced for IBM). Consequently, you won't find Ubuntu or Mandriva on its IdeaPad netbooks – the only option is SUSE Linux Enterprise Desktop. Lenovo shipped our review unit with Windows XP, but thankfully there was some quality Splashtop action to show the world how much Linux rocks, and straight away we installed Ubuntu Netbook Remix.

From a visual standpoint, the IdeaPad looks exactly like a normal, boxy business laptop that's been zapped with a shrinking ray. That's not to say it looks bad – far from it. But it doesn't have the curves and bends of the others on test here, and looks very serious from the off. We're not convinced about the usability aspects though. The keyboard is decently sized with full cursor keys, but it's quite rattly and the right-hand Shift key is the same size as the other keys, which isn't pleasant to work with (especially after the monster key on the Aspire One and MSI Wind.

Wrong way

Also, the Fn key is at the far bottom-left of the keyboard, where you'd normally expect Ctrl to be. It wouldn't be hard to swap the Ctrl and Fn keys around to provide a more familiar keyboard layout, as is the case on most of the others in this test. Then there's the trackpad: it's tiny. Really, really tiny. The buttons are crisp but very loud too. These aren't total usability calamities, but they do detract from what is otherwise a mostly well constructed machine, albeit with a slightly flexible glossy screen.

The IdeaPad is the only netbook in this test that includes an ExpressCard slot. This is on the right-hand side, along with headphone/microphone jacks, a USB port, an Ethernet socket and a Kensington lock (built neatly into the hinge). The left-hand side sports a reassuringly large air intake, plus VGA output, an SD card slot and a USB port. We're glad that the ventilation is so accommodating – the machine gets quite hot to the touch underneath when the CPU is heavily loaded.

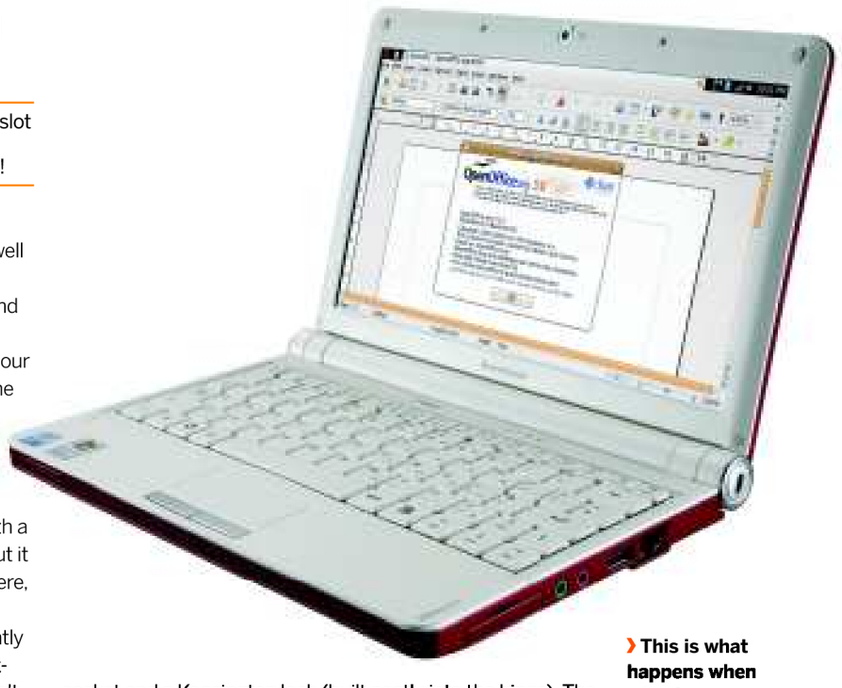

This is what happens when a business machine gets out of work, takes its specs off and loosens its tie. Maybe.

Tubular swell

Lenovo supplied us with the six-cell battery (£50–60) for our testing, which adds a tubular bump along the rear of the machine; imagine a couple of packs of Polos put end-to-end and you get the idea. The results in our benchmarks were impressive though: the battery lasted four hours and 16 minutes in our intensive test. For light usage in spates throughout a working day, you won't need to carry a charger with you.

"The battery lasted four hours and 16 minutes in our test."

Software-wise, Ubuntu Netbook Remix runs excellently, and the 160GB hard drive is a snappy performer compared with some of the sluggish SSDs we've seen. The speakers deliver good bass and point towards you from the front side of the unit's base. Overall, the Lenovo is a good machine, but there are a few usability niggles which could get on your wick.

Big bucks drives

SSD (Solid State Drives – aka flash drives) are definitely the future, and we can expect hard drives to be phased out gradually over the next few years. But large and well-performing drives are currently very expensive. We got hold of an OCZ Vertex 120GB SSD that sells around the £350 mark and installed it in the MSI Wind. Here's the bootup benchmark results:

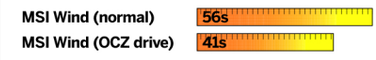

Not a huge difference then. Writing a 100MB file took two seconds, which is certainly at the top of the SSD performance but not a massive leap over hard drives. Until these devices fall dramatically in price, we wouldn't recommend them for netbooks – they're overkill.

Say hello to the brightest hard drive activity light in the netbook world. It's a bit distracting at times.

LINUX FORMAT Verdict

Performance	8/10
Usability	6/10
Build quality	7/10

A decent all-rounder, but please, Lenovo: swap those Fn and Ctrl keys, and give the trackpad more room!

Rating 7/10

MSI Wind U100

Price: £264 (www.scan.co.uk, three-cell battery version)
Web: http://uk.msi.com

Recommended for: Power users
Not recommended for: Rough environments

The MSI Wind U100 performed consistently strongly in our benchmark tests.

MSI (Micro-Star International) isn't a household name in computers: the Taiwanese firm sells most of its machines to other vendors, who add their own logos and branding. We've already seen this with the LG X 110 earlier in the test, but sometimes MSI likes to go it alone, and consequently we have the Wind series of netbooks. (There's also a series of desktop computers from MSI called Wind.)

Straight away there's one thing clear about the U100's casing: it's a total fingerprint magnet. That's nothing that a quick wipe with some cloth can't fix, but if you like your electronics looking perpetually shiny, you might find it frustrating when the machine catches a certain light and a smothering of smudges appear. The left-hand side of the unit houses an excellently large air intake vent, along with the Kensington lock slot, power socket and two USB ports. On the right you'll find another USB port, the SD/MMC card slot, headphone and microphone jacks, VGA out and Ethernet.

> **"UNR configured Bluetooth straight out of the box."**

Beneath the extremely rigid screen is a tough, non-rattly keyboard that, like the Lenovo IdeaPad, has the Fn key in the bottom-left corner rather than the Ctrl key, which takes some getting used to if you're a heavy Ctrl key user. The right-hand Shift key is to the left of the up cursor key, which is normal on desktops but not on netbooks. Otherwise it's a very good keyboard. The trackpad is sufficiently deep, with the buttons going almost to the edge – while it looks like one button in the photos, it holds two switches. We'd like MSI to add more width to the trackpad though; after all, there's plenty of room.

Power up

MSI supplied our Wind with a six-cell battery (available online at £35–40) that juts out of the bottom of the unit by 1cm and adds a bit to the weight, bringing it close to the Eee 1000 on the scales. However, we have some reservations about the build quality. It's not bad – it doesn't feel flimsy or as if it'll break apart in your hands. But the plastics creak with pressure in places and the hinge isn't as tight as we'd like. With a good case and a careful owner we can't see any big problems occurring here, but we have more confidence with in the main competitor in its class, the Asus Eee PC 1000, being able to take punishment.

Hardware recognition

The Wind is one of the few netbooks in this guide to include Bluetooth, and Ubuntu Netbook Remix configured it straight out of the box. The webcam was not detected, however. In terms of wireless signal pickup, boot time and overall performance, the Wind is mostly on a par with the Asus Eee PC, although the Eee's SSD drive is slower than the Wind's hard drive.

If you're looking for a larger-sized, almost-a-notebook netbook with a hefty battery, your two options in this guide are the Wind and the Eee. There isn't a huge amount to choose between them in terms of performance, but the Eee has a better keyboard, sturdier shell and slightly lower price point, so we recommend it. This is still a respectable showing from MSI though.

LINUX FORMAT Verdict

Performance	9/10
Usability	7/10
Build quality	6/10

» *A very good performer - not as noisy as the Asus Eee PC, but not quite as rigid either.*

Rating 9/10

Toshiba NB100 11R

Price: £229
Web: www.toshiba.co.uk

Recommended for: Business users who need max portability
Not recommended for: Anyone with big fingers

Now we're getting into tiny territory. When we first got our hands on the original Eee 701 we found the keyboard a bit cramped, though usable for hunt-and-peck typing. But as the netbook market started to expand and screens of around nine inches across became the norm, it was clear that adding just an inch of width on to a keyboard could make a major difference. And that's the situation we have here. The Toshiba NB100 sports the same type of keyboard as the Eee 701, albeit more noisy and rattly, and if you've got big hands you'll find it quite pokey in comparison to larger machines.

Still, this isn't to be seen as an error by Toshiba: the company is striving for portability here. Width-wise it's the second smallest (to the ONEt) netbook in our test, pipping the Dell Mini 9 by a centimetre. It's a bit taller than the Dell, though, largely due to the unusual battery that protrudes from the rear (more on that in a moment). Cosmetically, the NB100 is much like the IdeaPad – it really does look like a scaled-down laptop. However, it also has a glossy screen which looks lovely at some angles, and horribly reflective at others.

Toshiba has been more creative than the other participants with regards to port placement. The headphone and mic jacks are mounted on the front-facing side of the machine, while there's a single USB port and Kensington lock on the left-hand side, a couple of extra USB ports and the SD/MMC card slot on the right, with Ethernet, power and VGA on the rear. The build quality is first-rate, aside from the noisy keyboard – it's all very firm and sturdy.

That tiny keyboard makes the NB100 11R hard work for adult humans.

Now, about that battery...

The battery sticks out of the back by about a centimetre. Now, that's no big cause for concern, right? Other netbooks in this test have six-cell batteries that also bump out a bit! But the NB100's only has four cells, and the reason it sticks out is because it's not very wide. On the underside of the Aspire One, for instance, the battery stretches to around 90% of the width of the machine; on the NB100, it's closer to 60%. Toshiba's designers have clearly worked hard to cram everything in to a small shell, but we reckon a full-width battery that integrates fully with the rear would be better than a narrower battery that juts out.

> **"Toshiba has worked hard to cram everything in."**

Moving on to performance: there are a few problems here. The boot speed isn't great, and the wireless network pickup strength jumps around alarmingly even when the machine is static, leading to the weak results in the wireless table on p46. Overall performance felt choppy under Ubuntu Netbook Remix – video files stuttered and there were longer than usual delays in opening programs. It doesn't make the machine unusable, but it can be grating.

All things considered, the Toshiba NB100 11R has a hard time being competitive in this group test. It's very compact and has a serious, business-like facade, so if you just need something small to take to meetings and do a bit of office work, it's a solid and quiet little machine. But for most users we recommend sacrificing a bit of portability – just an extra inch in width – to improve usability.

LINUX FORMAT Verdict

Performance	5/10
Usability	5/10
Build quality	8/10

Workable as a business but let down by stuttering performance and a teensy keyboard.

Rating 6/10

Moblin and SUSE

Intel's Linux-based mobile OS, Moblin (see **LXF118**'s What on Earth), has received a healthy shot in the arm thanks to the work of Novell's SUSE team. As we were writing this article, Novell was preparing to demonstrate its "implementation of Moblin into the OpenSUSE codebase" as Holger Dyroff from the company's Business Development unit told us.

Right now, SUSE is supplied with a handful of netbooks from MSI and HP; with its enhanced version of Moblin Novell is hoping to reach out to a wider market, especially businesses. "Enterprises can see netbooks as thin-client replacements" said Dyroff, and the company has a "huge interest to make a business out of it".

Will it be the next big challenger to Ubuntu Netbook Remix? By the time you read this, Novell should have publicly released its work at **http://en.opensuse.org/Moblin** – give it a go.

Who's better? Who's best?

Almost all of these notebooks are based on the exact same Intel chips, and yet there's such a variety in size, weight, build quality, drive performance and price. Aside from the Elonex ONEt's low-end CPU and the Aspire One's lethargic SSD, performance doesn't vary enormously between them: all but the ONEt are perfectly capable of web browsing, YouTube, office work, and even a spot of on-the-go programming, providing you're not working on an outrageously demanding 3D showcase.

In each review we've looked at the machine from the point of view of the right kind of user for each model, so you've probably already got an idea of what would suit you best. If you're approaching this solely from a financial perspective, though, here's what we recommend:

» **Under £150** The Aspire One. Absolutely. Even if you see the ONEt (or another netbook based on the same internals) for under £99, avoid it. The money you save isn't worth the pain of sluggish web browsing with many sites out of bounds.

» **£150 to £250** If maximum portability is crucial, go for the Dell Mini 9. The keyboard is a stumbling block, but it's a very well designed and constructed machine. If you can tolerate something more bulky and noisy, get the Eee 1000 – a great all-round machine.

» **£250 upwards** We really love the LG X110; it's just a shame it's only available with Windows. If you're tempted, keep checking online in case LG expands its options and offers a Linux version, in which case it's a must-buy. Otherwise you should consider the slightly weaker, but still good, MSI Wind.

Of course, if you already own any of the netbooks we've tested here, we'd love to hear your opinions! Head over to our forums at **www.linuxformat.com/forums** and jot down your own experiences – have you fallen in love with a particular machine? Did the ONEt's rainbow of colours tempt you? Have you modded your Mini 9 to add Fx keys? Let us know! LXF

Boot time

Measured from pressing the power button to arriving at the desktop and disk activity stopping. All benchmarks are with UNR 9.04 except for the Elonex ONEt.

Acer Aspire One	78s
Asus Eee 1000	43s
Dell Mini 9	42s
Elonex ONEt	33s
Lenovo IdeaPad	54s
LG X110	56s
MSI Wind U100	56s
Toshiba NB100	63s

Battery life

Measured with power savings disabled, full screen brightness, connected via Wi-Fi and playing a video on loop from the local drive.

Acer Aspire One	1 hr 43 mins
Asus Eee 1000	3 hr 42 mins
Dell Mini 9	2 hr 52 mins
Elonex ONEt	3 hr 7 mins
Lenovo IdeaPad	4 hr 16 mins
LG X110	1 hr 32 mins
MSI Wind U100	3 hr 21 mins
Toshiba NB100	3 hr 5 mins

Drive write speed

Determined by copying a 100MB file locally and running **sync** immediately afterwards to flush the changes to the drive. Averaged from three trials.

Acer Aspire One	27s
Asus Eee 1000	8s
Dell Mini 9	6s
Elonex ONEt	73s
Lenovo IdeaPad	4s
LG X110	3s
MSI Wind U100	3s
Toshiba NB100	3s

Wireless pickup

Signal quality rating from *NetworkManager* status applet (**iwlist** command on the Elonex ONEt), access point six metres away, with brick wall in between.

Acer Aspire One	48
Asus Eee 1000	72
Dell Mini 9	40
Elonex ONEt	86
Lenovo IdeaPad	70
LG X110	68
MSI Wind U100	74
Toshiba NB100	38

Name	Acer Aspire One A110	Asus Eee PC 1000	Dell Mini 9	Elonex ONEt	Lenovo IdeaPad S10e	LG X110	MSI Wind U100	Toshiba NB100 11R
Price	£139	£249	£228	£119	£289 (three-cell battery)	£284	£264 (three-cell battery)	£229
Size	248x170x32 mm	266x190x38 mm	233x168x33 mm	213x143x32 mm	250x184x34 mm	261x181x35 mm	258x182x40 mm	224x190x37 mm
Weight*	910g	1290g	1000g	610g	1050g	1150g	1225g	1020g
Display	8.9", 1024x600	10", 1024x600	9", 1024x600	7", 800x480	10.1", 1024x576	10", 1024x600	10.2", 1024x600	8.9", 1024x600
Video	Intel 945GME	Intel 945GME	Intel 945GME	JzSOC chip	Intel 945GME	Intel 945GME	Intel 945GME	Intel 945GME
CPU	1.6GHz Intel Atom N270	1.6GHz Intel Atom N270	1.6GHz Intel Atom N270	400MHz XBurst	1.6GHz Intel Atom N270	1.6GHz Intel Atom N270	1.6 GHz Intel Atom N270	1.6 GHz Intel Atom N270
RAM	512MB	1GB	1GB	128MB	1GB	1GB	1GB	512MB
Storage	8GB SSD	40GB SSD	8GB SSD	2GB SSD	160GB HD	160GB HD	80GB HD	80GB HD
Battery	3 cell, 2200mAh	6 cell, 6600mAh	4 cell, 2200mAh	2 cell, 2100mAh	6 cell, 4360mAH	3 cell, 2200mAh	6 cell, 5200mAh	4 cell, 3800mAh
Ports	3xUSB, 1xSD, 1xSD/MMC, Ethernet, VGA	3xUSB, 1xSD/MMC, Ethernet, VGA	3xUSB, SD, Ethernet, VGA	3xUSB, SD, Ethernet	2xUSB, SD/MMC, VGA, Ethernet, ExpressCard/34	3xUSB, SD/MMC, Ethernet, VGA	3xUSB, SD/MMC, Ethernet, VGA	3xUSB, SD/MMC, VGA, Ethernet
Webcam	0.3MP	1.3MP	1.3MP	No	1.3MP	1.3MP	1.3MP	0.3MP
Wireless	Atheros AR242x	RaLink RT2860	Broadcom BCM4312	ZyDAS ZD1211B	Broadcom BCM4312	Realtek RTL8187SE	RaLink RT2860	Atheros AR242x
Bluetooth	No	Yes, not recognised by Ubuntu	No (£35 extra)	No	No**	No	Yes	No
OS	Linpus Linux Lite	Xandros Linux	Ubuntu 8.04	CELinux-based distro	SUSE Linux Enterprise Desktop	Windows	SUSE Linux Enterprise Desktop	Ubuntu 8.04

*Including battery **Only on Windows version


Donate your computers to charity
www.computeraid.org/donate
Free Blancco data destruction
WEEE compliance guaranteed
Zero % materials end up in landfill
Legal liabilities covered
Computer Aid International
Email: enquiries@computeraid.org
Tel 020 8361 5540
Registered charity no. 1069256

The Games Master

Michael Simms heads both Linux Game Publishing and Tux Games. That means you have him to thank for bringing *X3*, *Cold War* and more to your desktop.

LINUX FORMAT Interview

On the back of the dotcom boom, Michael Simms ploughed £350,000 of his own money into a games company with the intention of bringing some of the most playable Windows titles to Linux. Almost 10 years later, Linux Game Publishing, which specialises in porting Windows titles, is still going strong, releasing several titles every year. We caught up with Michael on a recent trip and asked him about where the company will go from here.

Linux Format: What inspired you to start making games Linux-compatible?
Michael Simms: I started using Linux when I was at university, so I've been doing it for a long time. I did a few jobs in the Unix field and got to hear of Loki Software, who had just decided to make *Civilisation: Call to Power*. I got on to the beta for that, but I found it was hard to buy a copy of it when it came out. So I contacted Loki about becoming a reseller and that's what started Tux Games. When it became obvious Loki was going under, it was like: 'crap, we're going to have nothing to sell'. So we went to a company we knew weren't able to make a deal with Loki, Creature Labs, and came to an agreement with them. We started off by publishing *Creatures 3* and went from there.

LXF: What do you think Loki did wrong?
MS: Loki overestimated the market. It would spend a lot licensing a triple-A title and not generate enough sales, but carry on doing that again and again. A classic example was its *Quake 3* special edition where it made 50,000 tin boxes and only sold a few thousand.

LXF: Didn't you do a similar thing with *X3*?
MS: That was a limited edition of 500 rather than 50,000! We did it slightly differently. Just 500, and we won't be making any more. We'll carry on making the standard edition until whenever, but try to avoid making the same mistakes as Loki.

LXF: After deciding to port a game to Linux, what's the next step for you at LGP?
MS: Once we've made the agreement, we get hold of the source code and then we just do whatever we need to do for the port. Usually, ports are fairly similar.

LXF: Do you choose games with a similar back-end?
MS: No, we choose games based on playability. I personally pick out a lot of the games because they're what I like! But we've also got a few other people that we trust to give a balanced view of things.

LXF: Is there a massive difference between taking on something like *X3* and a 2D puzzle game?
MS: We aim to do fifty-fifty top-end games to entry-level games so that we can pay equal attention to companies behind titles like *Jets'n'Guns* (see page 29 for our review). We concentrate on both to make sure that we're still seen to port big games, but we do small games so that smaller companies also have a route into the Linux market.

LXF: How long does a game like *Jets* take to port?
MS: Well, with *Jets*, we didn't actually do most of the port. Rake In Games did it instead, we just added some polish and work at the end ... To do a port of something like *Jets* would take one developer a couple of months. Maybe a bit less. *X3* – that's more a team of four developers for five to six months.

LXF: Does LGP pay their wage full-time?
MS: We have a few people on a salary, but most are on a commission basis.

LXF: How do you find people to work on a game?
MS: With great difficulty. When the game comes out, and people start getting their commissions, they usually end up with a wage that's roughly appropriate for the work they've done. But it's hard because they don't get money in advance.

LXF: And if there's a year's delay, as there was with *X3*, they don't get paid?
MS: Exactly. It's a bit of a problem. But we've got some good people on board now who are getting some royalties from previous games, so they're able to work on new games without worrying too much.

LXF: How far along are your own technologies?
MS: The multiplayer, from our perspective, is fairly mature now. We've got it in a number of games and it seems to be working well. We've released PenguinPlay, which is our multiplayer matching service that we're aiming to put in direct competition with GameSpy. It's still suffering from having a low number of users at the moment, but it will grow.

LXF: Have you ever thought about porting a popular framework to build around?
MS: We have thought about that, but it would be taking us along the same lines as *Wine*. To get a good level of efficiency while doing something like that would be difficult. The company that originally did the port for *Knights and Merchants* [Runesoft] tried to do it with something called the *dexter* library. It does the job, but it's terrible for efficiency, it really is. A game that runs on a 500MHz machine on Windows ends up taking over 1GHz on Linux because of the extra overhead of the middle layer. So it can make porting quicker, but you get a lower-quality game.

One thing we won't compromise is on quality. Every single game we've had was delayed in one way or another – I won't let a bad game go out.

LXF: Was that the reason for the *X3* delay?
MS: There were a few issues with *X3*. It turns out that some versions of the Nvidia driver didn't quite work the same way. Don't even get me talking about the ATI driver, but we have to support it. Getting it perfect on all of them is what took a bit of time. We also had a delay of three months where we had to find a bug in the rendering engine. In essence, the random number generator didn't work under Linux. And because it's a random number generator, it was hard to work out that it was going wrong – it's random by nature!

ON PORTING X3
"We ported everything except the graphics engine in a week."

LXF: Is there anything you can learn from the process?
MS: For the *X3* port, we ported everything except the graphics engine in a week. The game engine was similar to the *X2* engine, but the graphics engine was just so different. We were thinking, 'we'll be done in a couple of months', but we weren't. But all credit to my guys, they did a lot of hard work on that port.

LXF: Are you tempted to create an original game?
MS: We do have a couple of original game ideas, one of which is about half complete. It will be a fairly simple game to start off with – we're not a big-budget company. This is completely new, no one knows about this. It's very simple, it's based on Sudoku. But it's an entirely new take on the game. I can't go into details because it's still a few months away, but we're hoping we'll be able to get it out in the next six months.

LXF: Which are your favourites games?
MS: I'll always love *Majesty* – I've played it end-to-end about three or four times. I thought *Cold War* was brilliantly done and although I wasn't a fan of the gameplay in *Postal 2*, I loved the message that the company was trying to put out. Because you can play *Postal 2* in the most violent and graphic way, but you can also play it without hurting a single person. I don't know anyone who's played it like that, but I like that the people who made *Postal* are saying you can get through this game without any violence. LXF

› Michael picks a lot of the ports, so they're mostly what he'd like to play himself.

Mac OS X and Linux

Cousins under the skin

Juliet Kemp delves deep behind the gloss of OS X and finds that Macs and Linux have a lot in common.

Linux is fantastic, don't get us wrong, but Apple's kit can also be appealing – especially its range of laptops. But if the thought of getting to know OS X puts you off, fear not: there's a lot of common ground between Mac and Linux that isn't immediately obvious. For starters, you may already know that under the hood Mac OS X is BSD Unix-based. This means you can fire up a terminal window and get your fingers into the operating system via the command line in much the same way that you can with Linux. There's also an increasing amount of Linux software ported to Mac. We've taken a look into some of the similarities and differences between Macs and Linux to help if you're thinking of becoming a multi-OS household.

Our expert

Juliet Kemp
Despite a long and happy history with Linux, Juliet is unreasonably fond of her MacBook. She firmly believes that Macs and Linux can co-exist peacefully.

Getting started

How to navigate the world of Mac OS X using both Finder and the Dock.

Fire up a Mac OS X machine, take a look at the desktop and, aside from the layer of gloss, the first differences you're likely to notice coming from a Linux environment are the Dock and the menu bar. Let's deal with the latter first – you see, instead of the OS having a menu bar for system tasks and each application having its own individual bar within its window, in OS X there's a single menu bar at the top of the screen that alternates between the two.

The right-hand side of the bar contains the date, information on your battery (where applicable), a wireless button, and various other configurable options. Meanwhile, the left-hand side holds the application menu of whichever program has your focus, or Finder if none has. As you move between programs, this menu bar will change accordingly. The Apple menu, accessed via the little Apple logo on the far left, has a basic list of system operations and information. For more of these options, you need to open Finder, which enables you to navigate through files and your computer.

› The Mac desktop with a Finder window open, the Dock at the bottom and the all-purpose menu bar at the top of the screen.

Great minds think alike

The Spotlight tool on the far-right – the magnifying glass icon – enables you to search both filenames and the contents of files for anything you want, which makes it roughly equivalent to *Beagle* on Linux.

Finally, there's the Dock at the bottom (by default), where you can shove a selection of frequently used apps, files and directories for easy access. You can also set it to hide itself automatically so it's not in the way, which is particularly useful if you're on a laptop with more limited screen space.

If you try this out and like it, or like the idea and want to give it a go on Linux, there are a few Dock-a-likes available. *Avant Window Navigator* is probably the most popular, but there are other options available as well.

Applications

Porting apps to OS X is easy if you rely on the groundwork laid by others.

Applications in OS X are handled pretty differently to Linux. By default, they're kept in the **/Applications** directory, but you can place them anywhere you like, because each application is an **appname.app** directory that masquerades as a single file in Finder. If you want to look inside that directory, you'll need to go to the Terminal, type **cd /Applications/appname.app** and take it from there. The idea is that OS X applications should all be self-contained: all the application data is kept in that one file, rather than being scattered in various places about the system.

However, there is some app-related data that's kept elsewhere: there's **/Library/** for system-wide information and **~/Library** for individual data. In essence, if it's something that a user might want to change – themes, plugins and so on – then you'll find it in one of the **Library** locations. You can also get at the contents of the app folder using the Show Package Contents option in the right-click (or Ctrl+click if you're using a single-button mouse) context menu.

Starting up daemons is no more complicated than it is in Linux, but the process is slightly different. Some can be started from the appropriate System Preferences tab, or alternatively there's **StartupItems** and **launchd**. Most Unix-like daemons can be run via one of these, but they work differently, so you'll need to rewrite (or find) appropriate start scripts. In theory, you should use **StartupItems** in OS X 10.3 or earlier and **launchd** for 10.4 onwards, but **launchd** can be fussy, so you might have to mix the two.

Apple's *Safari* browser uses *WebKit*, a HTML engine based on the KDE project's own *KHTML*.

If you're running *Fink* or *MacPorts*, your downloaded apps and packages live in either **/sw** or **/opt/local**, respectively. The setup process for each should add the appropriate **bin** directories to your **$PATH**, and this all works much the same as in Linux. Check your **$PATH** setup with the command **echo $PATH** if you have trouble using any *Fink* or *MacPorts* apps (see page 53 for more about both).

Quick tip

You can set up multiple desktops using the Spaces application, which is accessed via the System Preferences app. Here you can choose how many desktops to run and set up shortcuts to access them.

The command line

Your Linux experience comes into play at the heart of OS X.

So far there's been more that's different than in common between our two OSes, but the similarities become much clearer as soon as you open up the *Terminal* program found in **/Applications/Utilities**. Type **ls** and you'll see a directory listing, just as you'd expect. Entering **pwd** will reveal that your home directory is at **/Users/username/** rather than **/home/username**, but that's just a cosmetic difference, especially since **~** refers to the home directory and **cd** with no argument will return you to it.

Try **echo $SHELL** to show that *Bash* is your default shell. **Which** works too, along with the standard shell built-ins and other basic CLI commands. **Locate** is also present, but note that to update the database, you'll need to run **/usr/libexec/locate.updatedb**. You'll need to do this as root and using **sudo**, because there's no separate root account on a Mac. You can set one up in the Accounts tab of the System Preferences application, which also enables you to set various other system config options. You won't need it all that often, though – often **sudo** will do just fine.

If you're yearning to know more, man pages are available exactly as they are on a Linux system. This can be particularly useful, because some BSD commands don't work in quite the same way as Linux commands. For example, **grep** on OS X doesn't have the **--exclude-dir** option, which is a nuisance. The Alt+. *Bash* shortcut to enter the final argument from the previous command line doesn't work on OS X, either, so you'll have to use Esc+. instead. In general, however, Ctrl does what you'd expect it to on the command line, so you can still use Ctrl+A and Ctrl+E to jump to the start or end of a line.

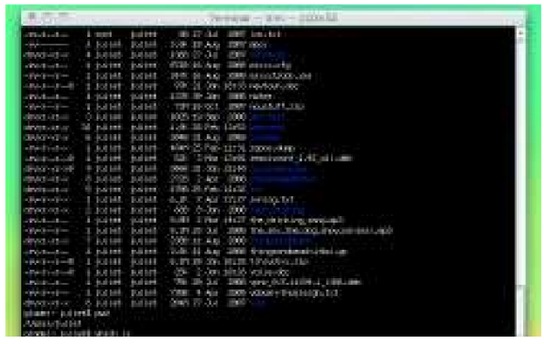

OS X's *Terminal* looks pretty much like any other console.

However, this can be slightly confusing given that in graphical apps, the Apple key takes the place of the Ctrl key on Linux. So, for instance, you'd use Apple+W rather than Ctrl+W to close a window. This is probably to do with the fact that *Terminal* is more Unix-like than the Aqua environment that lives on top of it and handles all the graphical jobs. When it comes to using *X11*, which is available for Mac and runs well in OS X 10.5, you'll just have to experiment a bit.

One final neat feature of *Terminal* is the **open** command. Type **open filename** and OS X will fire up the graphical program that's the default handler for that type of file, which saves a lot of time looking for the correct applications.

Quick tip

You can set a default handler via Finder by selecting a file of a particular type and hitting Apple+I. Set the Open With value, and then click Change All to make this the default.

GUIs and eye candy

You knew it was coming, so here it is: OS X is exceptionally pretty.

One of the features that delights Mac users is Exposé, which enables you to hit a button (or move the mouse to a specified part of the screen) and have all your open windows shown at once. There are clones available for Linux, such as *Komposé*, *Skippy* and a *Compiz Fusion* plugin called *Scale*, which may be worth investigating if you find yourself missing the power of Exposé on your Linux box. However, they're still not quite as shiny.

Some people find the Mac OS X Aqua look and feel is particularly attractive too. You can get Aqua-style themes for Gnome, although the extent to which these make it look like a Mac is a matter of debate.

A really neat GUI trick in Macs is the ability to preview files from Finder. Select the file, hit the Spacebar and a little window pops up with a preview of the file. You can then scroll through the directory, or return to the normal file listing.

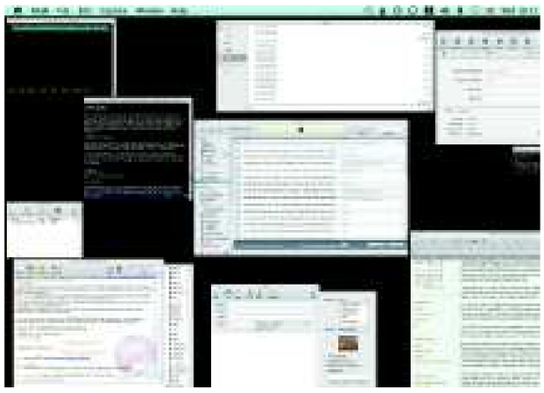

› **Exposé makes picking an app from those you have active easy. It's handy if you're often flipping between tasks.**

Filesystems

Getting to grips with the intricacies of HFS Plus and Time Machine.

The preferred filesystem in OS X is *HFS Plus*. This has Unix-style file permissions, journalling and extensive file metadata. By default, it's case insensitive but case preserving, although you can choose to format drives to be case sensitive instead. The default settings can cause problems if you're bringing files over from Linux, because ext2 and ext3 are case sensitive. It's particularly troublesome if you have files that share the same name and are only differentiated by case.

OS X can handle the various FAT filesystems and NTFS, but sadly won't deal with ext2 or ext3.

> **"Time Machine works similarly to Bacula, but it's plug and play."**

Samba, baby

However, what does work well is NFS and *Samba* mounting. You should know, though, that as a BSD-based system OS X will use a port numbered above 1024, whereas Linux NFS expects clients to use a privileged network port, which will be less than 1024.

There are two ways of solving this. The first is to add the **insecure** parameter to the **/etc/exports** file on the server:

```
/local    10.1.0.*          rw,insecure
```

If you prefer, and you have only one or two Macs on your network, you could do this to open the access only for those machines in particular:

```
/local    10.1.0.17         rw,insecure
/local    10.1.0.*          rw
```

The other option is to fix it at the client end, which is probably neater if you're the only one using the Mac. To do this, add **-P** to the command line via the *Terminal*.

```
mount -P server:/local /mnt
```

The other feature of Mac OS X that's worth mentioning while we're discussing filesystems is *Time Machine* – the archive and backup system built straight into OS X 10.5 (Leopard). This works like *Bacula* or any other similar incremental system, backing up changes as often as you define. However, the advantage of the program over similar offerings is that it's plug and play – once you've set up an external drive as your backup destination, *Time Machine* will automatically kick off whenever you plug it in. Not only that but the program will then carry on with the scheduled backups until the drive is full or unplugged.

It's probably possible to cobble together something to do this on Linux (using **udev** and *Bacula* or a similar incremental system), but on a Mac it's just there, ready and waiting. In short, *Time Machine* is absolutely fantastic for the backup-obsessed and a great way to help anyone who isn't make sure their backups happen on a regular basis anyway.

› ***Time Machine* is OS X's built in plug-and-play incremental backup interface. Mac users have no excuse not to back up!**

Mac on Linux and Linux on Mac

How to fuse the best of both and make a Frankenstein-style über OS.

Since Apple started using Intel processors, it's become more feasible than ever to run OS X on non-Apple hardware. Head on over to **www.osx86project.org** and you'll find a community who are working on doing just that. If you want to join them, be prepared for a fair amount of hacking, though: OS X out of the box only provides drivers for Apple hardware. The Kalyway install disk fixes this, but you'll probably still need to mess with the BIOS settings. Note that OS X is not free software, so you'll need a licence to run it.

On the flip side, you've been able to run Linux on Mac hardware for ages – there's a PowerPC port of Debian, for example – but it's even easier now that Intel chips are in use.

> **"It's now even easier to run a Linux distro on Mac hardware."**

You should be able to boot a live CD without any hacking – so it's easy to give it a go, or to use this once in a while if you just want to run a few Linux apps or check something out.

Finally, there's the virtual machine option. Again, if you want to run OS X inside a virtual machine on a Linux box, you'll need to pay for a licence. However, running Linux inside a Mac on a virtual machine is straightforward and painless.

Somefink special

Given that most Linux software is portable to Mac, though, there seems little need to bother with running a distro in a virtual machine. Since OS X is POSIX-compliant, meaning it conforms to the POSIX standards set by the IEEE to define Unix compatibility, software written for Linux or BSD can be easily recompiled to run on a Mac.

From our point of view, this means that we have access to a far bigger stack of software than just the programs natively released for Mac, even though that's an increasingly large quantity of software itself. There are two major projects that aim to make this process easier:

- *Fink* (**www.finkproject.org**)
- *MacPorts* (**www.macports.org**)

Courtesy of Apple

Both aim to provide OS X ports of pieces of open source software, in particular those that come from the Unix world. They're also fairly easy to use once installed. However, each one has its issues, and the chances are that you'll be a little further behind the trend than if you were installing the software from a Linux repository. In other words, the *Fink* stable package will be an older one than the package you'd find on, say, Fedora.

In terms of the differences between the two, *Fink* has a GUI available and can be used from the *Terminal*, whereas *MacPorts* is *Terminal*-only. In addition, *MacPorts* is built by Mac employees, which may make it slightly more reliable, and it tends to be a little more up to date. We recommend that you pick one and stick with it, because although they can run side by side, you're more likely to run into problems that way.

Making the most of Macs

As we've seen, there are a lot of similarities between Macs and Linux once you get into the OS a bit, which stems from OS X's BSD heritage. This means that it's straightforward to run both Macs and Linux boxes on your network and to switch regularly between the two. For laptops in particular, the 'it just works' factor of Macs is fairly compelling. There's also some great software available for Mac that isn't available elsewhere, especially graphical and music-making apps, but using it doesn't mean you have to skimp on using Unix programs.

The downside is that Mac hardware is costly and the OS isn't free. You probably don't want to shift over to Mac altogether, but it's well worth investigating a little, or at the least looking into some of the neat Linux software inspired by Mac design options. LXF

Next month
We'll show you how to put your Linux knowledge to good use in Sun's Unix-based Solaris.

Mac, BSD and Linux

Mac OS X was first released in 2000. Unlike previous Mac OSes, it was based on Unix, with code derived from Nextstep, FreeBSD and various other free software projects.

The underlying code of OS X is the Darwin operating system, which has been released under the Apple Public Source Licence. The version of Darwin in use since 2003 even meets the Free Software Foundation's definition of 'free software'. Darwin isn't exactly a BSD fork, although it contains a lot of BSD-derived code – check out the top line of the **man** output from your *Terminal* window.

However, Darwin isn't a fully usable setup by itself – you need various other components on top of it. This is why Mac OS X isn't free: on top of the free Darwin, Apple put various components (notably the Aqua interface that gives OS X its visual identity and Finder) to end up with OS X. The OS also comes with a bundled version of *X11*, although it's only installed optionally, not by default. To install it, go to the **Extras/XCode** sections of the installation DVD.

Importantly, OS X is POSIX-compliant, which means that a significant quantity of BSD/Linux software can be recompiled to run on it – hence the existence of projects such as *Fink* and *MacPorts* (explained *above*).

However, what you can't do with OS X installed on a shiny new Mac is fiddle around with the internals of the kernel or other system parts, which you can do with Linux. OS X is definitely Unixy, but it's not free software, so there's a limit to how much hacking you can do.

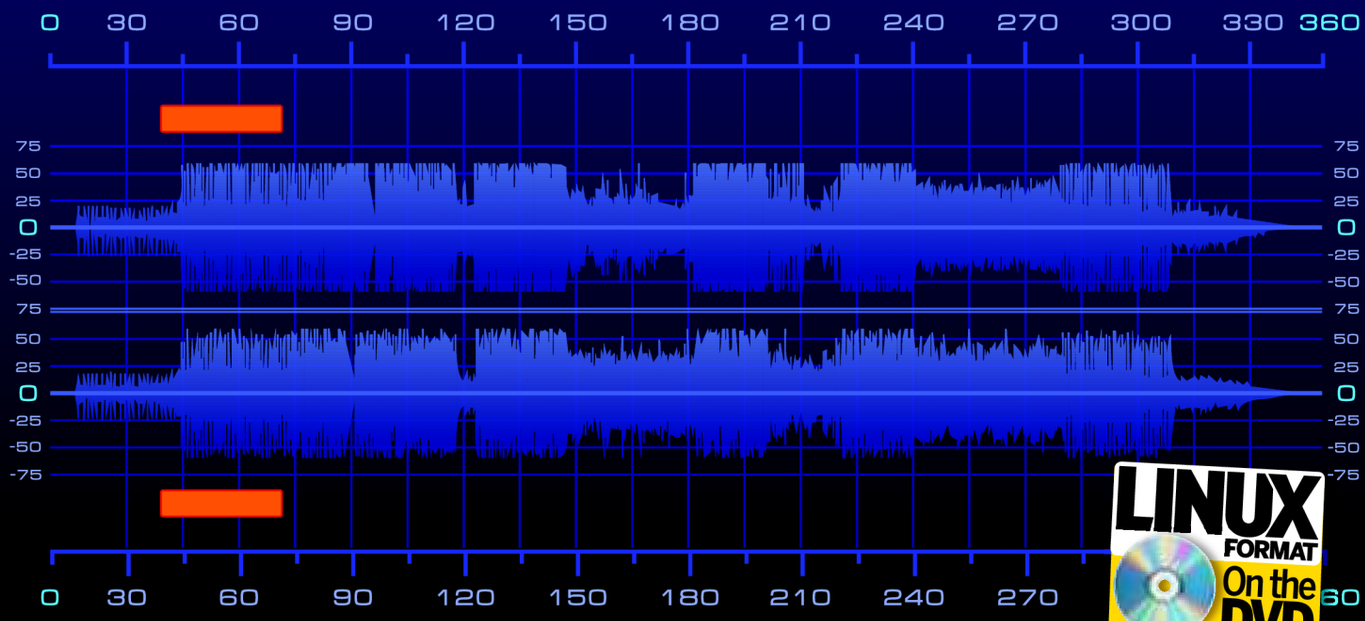

LINUX FORMAT On the DVD
- Sox 14.2.0
- Sweep 0.9.3
- Audacity 1.2.6

Sox: Hack your audio files

Our expert

Shashank Sharma is co-author of *Beginning Fedora* and has written about free software for over four years. He's also contributed to **Linux.com**.

Shashank Sharma shows you how to perform surgery on your audio files with Sox, the versatile command line sound processor.

As silly as it might seem, I first used *Sound Exchange* (*Sox*) to chop off a small chunk from a movie soundtrack so we could use it as a ringtone. There's a lot more to *Sox* than just hacking bits off an audio file, though. At its most basic, you can use it as a music player or to convert file formats of your audio files, but *Sox* has a bevy of more complex features to command as well.

Sox is installed by default on most current distributions, but yours may not have the latest version of it. The current release is 14.2.0, and if your version differs, your first step should be to remove the default version of *Sox* and install it again. You can remove *Sox* using your distro's graphical software management tool or use *Yum* and *apt-get*, if you're a command line junkie. Fedora users can remove Sox with

```
su -c "yum remove Sox"
```

Next, head on to the project's Sourceforge page at **http://sox.sourceforge.net** and grab the latest tarball. That's right, no pre-compiled binaries here people. Now run the command:

```
tar zxvf Sox-14.2.0.tar.gz
```

to extract the **Sox-14.2.0** directory from the tarball.

If you want Sox to support MP3, install the *libmad* and *Lame* libraries before installing *Sox*. While you're at it, also install *libmad-devel* and *lame-devel*. These are part of the software repositories for most distros, so you shouldn't have to hunt the web for them. Once everything is installed, browse into the freshly extracted **Sox-14.2.0** directory and then run **./configure** followed by **make** and **make install**.

You now have the latest version of *Sox* on your machine, with MP3 support, and no one's frowning. Well, maybe a few OGG fans. To appease them, you can convert your MP3 collection to OGG format with the command

```
Sox soundtrack.mp3 appeased.ogg
```

You can play either of these files with **play appeased.ogg**.

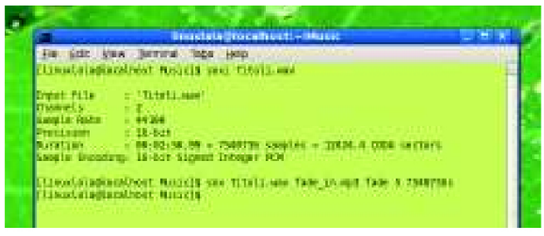

When not specified, the fade-out length defaults to the same value as fade-in length.

Performing surgery

Now it's time to start getting more complex – and what better way could there be to ease you in than making a ringtone? We'll be using the **trim** option to remove bits from the specified audio file, but first find a file to hack and then use **Soxi** on it. This command reads the header of the specified file and displays the exciting bits on the screen. When used

without any command options, **Soxi** displays all the information it can find in the header. However, you can control the information it spits out by using the various command options. The command **Soxi -d Titoli.wav**, for example, will only display the duration of the file **Titoli.wav** in HH:MM:SS. fraction format, while the command **Soxi -r Titoli.wav** displays the sample rate.

```
[linuxlala@localhost ~]$ Soxi Music/Titoli.wav

Input File      : 'Titoli.wav'
Channels        : 2
Sample Rate     : 44100
Precision       : 16-bit
Duration        : 00:02:50.99 = 7540736 samples =
12824.4 CDDA sectors
Sample Encoding : 16-bit Signed Integer PCM
```

For the trimming function we only need to know the Duration, be it in HH:MM:SS format or the number of samples. If you're only interested in the first 40 seconds of the file, run the following command:

```
Sox Titoli.wav 40_seconds.wav trim 0 40
```

The **trim** option requires two variables: start and length. The **0** and **40** in the command imply, respectively, the start position and length. In this case, we're taking 40 seconds from the start. If you want a section from the middle of the file, replace **0** with the relevant start position and specify the appropriate length. For example, the **middle.wav** file created with the command:

```
Sox Titoli.wav middle.wav trim 130 150
```

is a 20-second clip without the start and the end of the original file.

Fading out

There are two kinds of music lovers. Those who love fade-in and fade-out and those who don't – if you belong to the former camp, you can add a fade-in effect with:

```
Sox song.mp3 fade_in.mp3 fade 5 170
```

Here, **5** is the fade-in length and **170** is the stop time (the duration of the song) in seconds. You can apply both fade-in and fade-out as a single command if you keep to the following syntax:

```
Sox song.mp3 fade_in_out.mp3 fade fade-in-length [stop-time [fade-out-length]]
```

So, to create a five-second fade-in effect and a 10 second fade-out effect, use the command:

```
Sox song.mp3 fade_in_out.mp3 fade 5 170 10
```

Fade-out can't be added to the file unless you specify the stop time, but you can leave this out if you just want to add a fade-in effect.

> **"With Sox, you can easily add echo effects to your sound files."**

With *Sox*, you can easily add echo effects to audio files. When working with echo, you need to specify four values to the **Sox** command: gain-in, gain-out, delay and decay. Delay is the time difference between the original sound and its echo, which is also called reflection, while decay defines how loud the reflected sound is relative to gain-in. Delay and decay are specified in milliseconds (ms) and the gain-in and gain-out values describe the volume level.

Make mine echo

In addition to the basic echo, *Sox* also provides the **echos** function. Here, the original sound is used to create the first reflection, this reflection and the original are used to create the second reflection and so on. The command structure is almost identical to a regular **echo** command: you just give more delay/decay pairs.

Now, before you echo your audio files beyond recognition, remember the **play** tool we mentioned earlier? You can use this command:

```
play Titoli.mp3 echos 0.8 0.7 500 0.5 600 0.8
```

to figure out what all the number crunching does to your file, without altering the file.

We've only really scratched the surface of what *Sox* can do here, but we recommend you dig deeper and maybe explore *Sox*'s graphical counterpart, *Audacity*, too. LXF

Samples as duration

If you're uncomfortable with the HH:MM:SS.fraction time format, you can use the number of samples to specify the stop time. The **Titoli.wav** file for instance has 7540736 samples. We specify this in the command like so:

```
Sox Titoli.wav fade_in.mp3 fade 5 7540736s
```

The **s** at the end of the samples count is important because it tells *Sox* that we've specified the length in number of samples and not the end time of the file. Note that you can use the **-s** command option with **Soxi** to determine the number of samples in your file. To go back, divide the number of samples with the sample rate of the file to get the duration of the track in seconds.

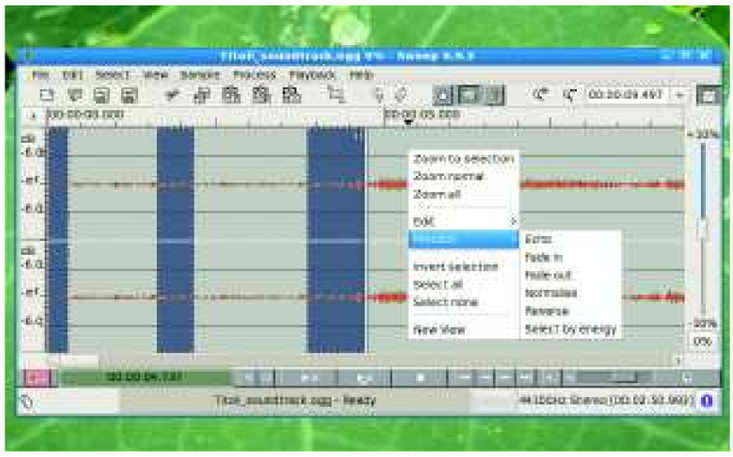

Another graphical alternative, to *Sox*, *Sweep*'s GUI allows for complex audio operations to be performed on different parts of the audio simultaneously.

Quick tip

While most of *Sox*'s functionality is invoked via the **Sox** command itself, the **play** and **rec** commands can be invoked separately.

A far more complex audio editor, *Audacity* allows multi-track mixing and many many editing effects, from a slick GUI.

Benchmarking and profiling

Our expert

Juliet Kemp is very lazy when it comes to optimising code, so she prefers to direct her efforts. That's why she likes profiling.

Sometimes your code is annoyingly slow. **Juliet Kemp** shows you how to identify why using the dark art of benchmarking.

Whether you're dealing with hardware or software, speed is likely to become an issue at some point along the line. That's a big topic to bite off, so here we'll chew specifically over when your code runs too slowly and what you can do to speed it up. This ranges from a rewrite to forking out for new kit, but first of all you need to find a baseline and work out where the bottlenecks lie. This is where benchmarking comes in – and it's well worth doing it properly before you start fiddling. The process is also occasionally known as profiling when you're looking at various sections of code, which we will be later on.

It's important to treat benchmarking as part of the process of code improvement. First write your code and ensure it does what you want to – only then decide if it does that fast enough. After all, there's little point in working to increase the speed if it's adequate as it is: your energies would be better expended elsewhere.

An ongoing process

As and when your code becomes sluggish, that's when you need to start benchmarking. After that, you can use the results (along with analysis and profiling) to establish where your effort will be best be spent. Once you've worked there, it's time to benchmark again and see if you've improved things enough. If not, pick the next area and keep going.

The important point is to make sure that you've got hard numbers (so you know what you've achieved) and focus your effort where it's going to have the most effect. There's no point in spending two days tinkering with your algorithm if the bottleneck is really the point when you're writing to the disk.

Initial benchmarking

The first thing to do is to benchmark the code as it currently stands. You'll almost always want to run the code multiple times in a row to get a representative result. This enables you

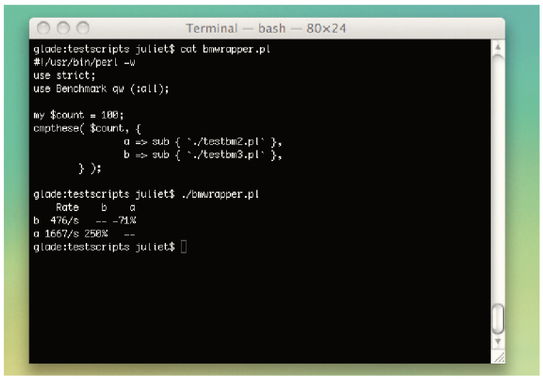

The output from a wrapper script using *Benchmark.pm*.

to average out the length of each run, which will fluctuate depending on what else the processor is doing at the time. The best way to do this is to use a wrapper script, such as:

```
#!/usr/bin/perl -w
use strict;
use Benchmark qw (:timethis);
my $count = 10;
timethis( $count, sub { `/path/to/myprogram` } );
```

This was written in Perl, but since it uses backticks(`) to call the code being timed, you can use it to time code of any language that's executable from the command line. Looking at our script, you'll see **$count** sets the number of iterations. An alternative way of doing this is to use a minus number to specify the minimum number of seconds of CPU time to run for – entering **-5**, for instance, will run for a minimum of 5 seconds of CPU time

The *Benchmark.pm* Perl module used above is also handy for benchmarking, even for non-Perl code (as before, you can just use backticks to call your code from within the Perl script). It should come by default with a reasonably up-to-date Perl install, but it can be installed via *CPAN* as well. Later, once you've improved your code, you can also use *Benchmark.pm* to tell you how much faster it is with:

```
#!/usr/bin/perl -w
use strict;
use Benchmark qw (:all);
cmpthese( 10, {
        a => sub { `/path/to/oldcode` },
        b => sub { `/path/to/newcode` },
    } );
```

Note that **cmpthese** will compare the number of code runs per second, rather than the overall time. However, this can be more useful to know than the absolute time, because it can be affected by CPU scheduling and other system features. The wrapper script also suppresses any output from the compared pieces of code to **stdout**. By default, **cmpthese** will time to the nearest integer second, but you can enable high-resolution wallclock timing using:

```
use Benchmark qw (:hireswallclock);
```

to give you timing to the microsecond instead.

Keep a record

These wrappers will output the total time taken to the console, but you might want to pipe the output into a file (use **wrapper.pl >> benchmark.txt**, which will add each run's output to that file rather than overwriting it) so that you have a record of your progress.

Note that you should always keep a copy of the old code around when you start editing it. This enables you to run comparisons (as above) and it gives you a check for what results you should be getting, because it's always possible to introduce bugs while you're optimising. It also acts as a safety net in case you screw up badly while trying to optimise. Generally, you should be using a version control system and committing new versions regularly as you edit: it's well worth the effort for the advantage of being able to roll back to a known good version should something go wrong.

OK, so it's time to run the first set of basic benchmarks. Remember not to do anything else while the benchmarking is running, or you'll mess with the results. If it's practical, you should also take several runs of your wrapper script, which itself is generating an average, and average the result to give you the most accurate value. Make sure you're not downloading something that might affect your processing or I/O speed as well. You want to do as many runs as you have the time or patience for – the more runs you do, the more accurate the result will be.

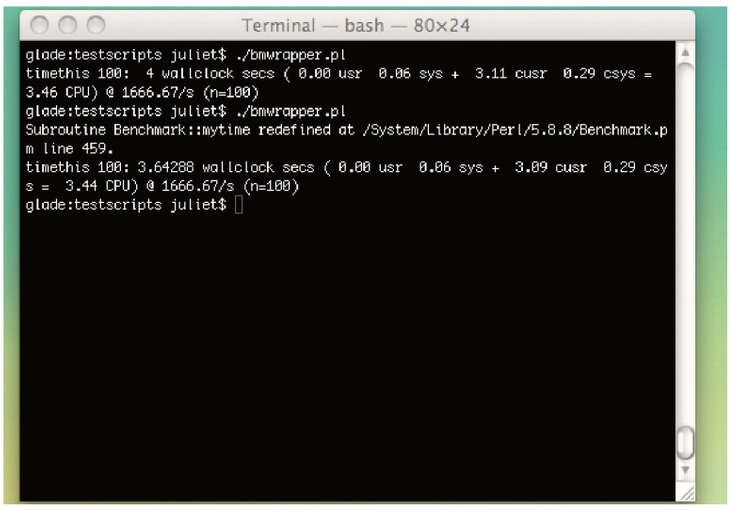

The first output we looked at has regular *Benchmark* timings, while the second uses hireswallclock. The difference between the two lies with the units of measurement – hireswallclock measures in microseconds instead of seconds.

Profiling: breaking the code down

Once you've got that overall value, the next thing is to fiddle with the code a bit to get some breakdowns. You'll need to look at your own code to see which sections you want to split out. Some suggestions are:

- Disk I/O sections (for instance, when you're reading from disk or writing to disk).
- Setting up and/or populating data structures.
- Algorithms and calculations.
- Anything that happens over the network (but you may not be able to do anything about this kind of speed problem).

For a compiled language, you'll probably have to do some rewriting or breaking up of your code at this point. Cutting out output to disk is straightforward and may just need a comment or two, but cutting out input from disk is harder because you'll need data for the algorithm to work on in order to get any results. Instead, you can cut out everything but the reading in of data, and subtract the average of those runs from your full runs average to give the average timing for the algorithm.

> **"Later, Benchmark.pm can tell you how much faster your code is."**

To see an example of that, let's say there are three main sections to your code: reading in data, running calculations on the data and writing out the new data. Create three versions of your code:

A: The full version.
B: One for reading the data in (no calculations or output).
C: Finally, another for reading the data in and running the calculations (no output).

Benchmarking these versions gives you the information that A takes 4 seconds, B takes 1 second, and C takes 3.5 seconds. That means that reading the data in takes 1 second (from B). Outputting the data must take 0.5 seconds, because that's the difference between A and C. The calculations, therefore, take 2.5 seconds. You can calculate »

Quick tip

Even if you don't have any problems with your code's speed currently, it's worth running a couple of benchmarking tests to ascertain the overall value. Then if you think it's running slower later, you can check out whether it really is, or whether you're just getting less patient.

this using the formula: A - B - data output = run time (in our case, 4 - 1 - 0.5 = 2.5). That's the time taken for a full run, without the time taken for data input and data output.

You can also use this technique to get timings for other parts of your code that are hard to separate out.

An alternative is to put in print lines at relevant points, outputting the clock time at that point. The Perl script below will generate one line per code run:

```
#!/usr/bin/perl -w
use strict;
use Time::HiRes;
sub time_print;
print time_print . ":";
# code here
print time_print . ":";
# more code
print time_print . "\n";

sub printtime {
        my ($t1, $t2) = Time::HiRes::gettimeofday;
    my $time = "$t1." . sprintf("%05d", $t2/10);
    return "$time";
}
```

And it will give you a readout like this:

```
1240905933.05204:1240905934.05249:1240905935.05264:124
0905936.05312
```

Each value in the pairs here is separated by a full stop and each of the four pairs is separated by a colon. In a pair, the two values are in seconds since epoch and microseconds; the first value is at the start of the script and the last at the end. Using the *Time::HiRes* module enables you to get a count in microseconds rather than just integer seconds.

> **"It's important to locate where the bottlenecks are in your code."**

Handling the data

Run the code multiple times, pipe the output into a data file, and then use this Perl script to handle it:

```
#!/usr/bin/perl -w
use strict;

my $datafile = "testout.txt";
my @timearray;
my $count = 0;

open DATA, $datafile;
while (<DATA>) {
    my @time = split /:/;
    push @{ $timearray[ $count ] }, @time;
    $count++;
}
close DATA;

my @result;
for my $rowref ( @timearray ) {
    my @row = @$rowref;
    for my $i ( 0 .. ( $#row-1 ) ) {
        $result[ $i ] += $row[ $i+1 ] - $row[ $i ];
    }
}

for my $i ( 0 .. $#result ) {
    print "Section " . ($i+1) . " average = " . $result[ $i ] / $count
. "\n";
}
```

This will print output (*as shown in the screenshot, below left*), but note that the test script has only been run three times, which isn't enough for adequate profiling.

After all this, you should now have an idea of how long your code takes to run as an overall package and at specific points. Next stage: speeding it up.

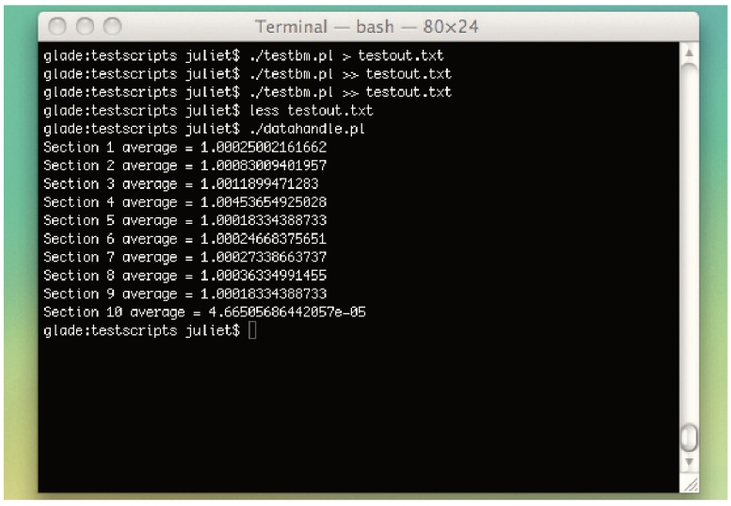

Running the output parser to do some basic profiling on sections of code.

Calculating timings

A further alternative to consider when calculating timings if you're using Perl and *Benchmark*, is to find differences between the times using:

```
my $t1 = new Benchmark;
# code
my $t2 = new Benchmark;
my $td = timediff($t1, $t0);
print "First section took $td \n";
```

However, the output from this is less easy to parse.

Resolving issues

As we discussed above, it's important to locate where the bottlenecks are in your code, so you can focus your effort where it's likely to be most use. You may well be able to see where the problem sections are just by eyeballing the figures, but if you can't, try using graphing software such as *OOo*, *KChart*, *Guppi*, or *Graphviz*. Alternatively, you can use **gnuplot** if you prefer to work at the command line.

If your code is I/O heavy, there's a fair chance that that'll be your bottleneck. Possible speedup solutions include:

- Ensure that you're working locally, rather than over the network. You can always copy the files on to a remote drive once you're done with them.
- Buy new kit. It's possible that your old kit just isn't up to the job any more. You can always retire it to somewhere where speed is less important.
- Batch up I/O jobs so that they're all managed at once. In particular, if you're reading in data, you want to read it once and then run all your processing from memory. If you're going to run this code multiple times, consider setting up a wrapper to read the data in once at the start.
- Lower the niceness of the process.

You can use **hdparm** to have a look at your current disk parameters (for example, **hdparm -v /dev/hda**). A straightforward change that may increase disk speed is to turn on Direct Memory Access (**hdparm -d1 /dev/hda**), you should then benchmark again to see if it's worked. You can also change the I/O support value with the **-c3** option, which may improve things slightly. Experiment with other **hdparm** values, but be warned: some of them can be dangerous. Read the man page before going ahead. If you want to keep the changes you make, run **hdparm -k /dev/hda**.

Filesystem tweaks

You can improve file access times by setting the filesystem to avoid making a record every time a file is accessed – thus removing a write to the directory for every file opening. To do this, edit **/etc/fstab** to add **,noatime** to the list of options in the fourth column for any drives you wish to change. Afterwards, remount everything with **mount -a**. The downside to this is that you no longer have a record of file access. Be aware that changing this may cause problems for some mail and backup programs. You could set up a different partition just for your code data and only set **noatime** to deal with that partition, though.

If you're considering buying new kit, see if you can borrow time on someone else's better-specced machine to run some benchmarking before you spend the money. It's better to know up front what kind of result you're likely to get.

If your bottleneck is with your algorithm, there are plenty of resources for analysing algorithms. For example, there's lots of analysis available of sorting algorithms, which enables you to choose the algorithm that performs best for your data type – in general, a binary tree sort will be a better bet than a bubble sort, for example. Going into this further is beyond the scope of this article – there's plenty of research information available online.

Do bear in mind that optimising algorithms can be extremely hard work, so it pays to be really sure that you want to spend your time here for what may be a fairly limited speedup. Having said that, sometimes the right choice of algorithm can make a massive difference.

Strace and ltrace

Strace will output a trace of system calls made by a program. This will generally be overkill for benchmarking, but there are a couple of options that may be useful. For example, **-c** records the time for each system call, while **-r** prints a relative timestamp for each system call and **-t** instead prints the absolute timestamp (use **-tt** to include microseconds). **-T** shows the time spent in system calls. You can use the **-e** option to trace only a particular set of system calls.

ltrace does much the same thing for library calls. Again, **-c** will count the time for the calls and return a summary on exit, and **-t** or **-tt** will show the time of day at the start of each line.

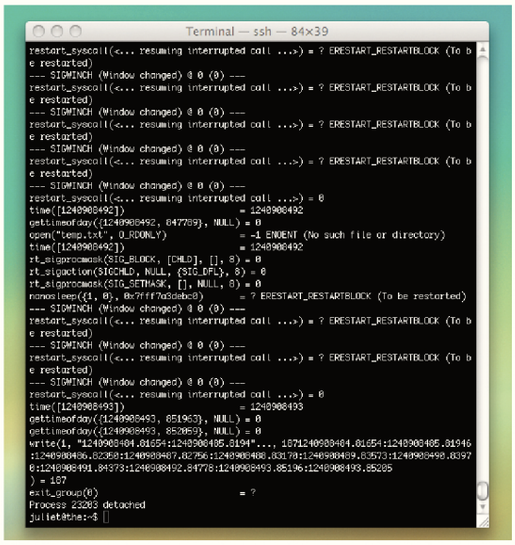

Strace running on one of the sample Perl scripts.

Parallel processing

If your algorithm can't be improved, another option is to look at using parallel processing to run your software faster. You'll probably need to look into writing software specifically to take advantage of parallel processing: in effect, you need to be able to split your analysis into usefully parallel chunks. For example, if you're doing data analysis on a very big corpus of data, but not all of that data interacts with each other, try running the same algorithm in parallel on different machines, with different chunks of data. There's not room in this article to go into the details on this, but there are tutorials available online to get you started. Unfortunately, not all problems are parallelisable!

Quick tip

The **relatime** option is an improvement on **noatime** – if you're using Ubuntu, you should try it out.

Wrapping up

To get more detailed information on where your code is slow if you're writing Perl, you can try running *Devel::NYTProf*, or its predecessor *Devel::DProf*, which are both nuts-and-bolts benchmarkers with good documentation. Using these tools, you may be able to find places where particular modules are causing problems and think of ways to take them out. Remember to run the profiling again afterwards to check that you've achieved something.

If you're using Perl or another scripting language and you can't find any way of speeding things up, but your speed really isn't acceptable, it may be time to rewrite in a compiled language such as C, C++ or Java. However, before you dive in and begin to radically change your code, you ought to be totally sure that this is worth your time and all the extra effort. Have another look at your benchmarking results and make sure you've done any other speedup work first to see if you can avoid the hassle.

Once you've finished optimising the code section that you chose, it's really important to make sure that your new code still works the same way as your old code. Unit tests are the best way to go about this, and they provide another good reason to keep your old code around, to check the two against each other. It's very easy to introduce new bugs when optimising. Keep a note of any changes you make so you don't revert them again later, or forget what you did when you want to do something similar again.

The other thing you must do at this point is benchmark again! Find out how much time you've saved overall, bearing in mind how much work time it cost you, and think about whether that's enough. If you need to lose a few more seconds, is there anywhere it's feasible to make that happen? Or do you need to rethink the code altogether?

Benchmarking and the code improvement process can be great fun if you go about it with a clear aim in mind, and are aware of what you can and can't improve. However, it can be an exercise in frustration if you do it without thinking it through clearly first. Make sure you're firmly in the former camp first, then have fun reworking your algorithms and overspending on a selection of shiny new kit. LXF

What on Earth is APML?

Marco Fioretti explains how you can let websites know what interests you.

Huzzah – another acronym for me to forget! What does this one stand for?
Attention Profiling Markup Language, which enables websites to keep track of what's caught your interest and help them better understand your needs.

Is this about online ads? They don't stress me at all – I filter them out.
Filtering out all web advertising is a sure-fire way to terminate many free, independent websites that provide content you couldn't find anywhere else. Are you sure you want that? Also, advertising is only one place where APML could make a big difference.

Hmm... OK, so how is APML supposed to improve things?
Imagine the internet is like walking into a crowded flea market with all the vendors jumping in front of you and yelling: "Look at me, buy my merchandise, buy it, buy it, buy it!"

It's not difficult, that's why I tend to filter them out...
Sure, but now imagine you could carry a sign saying, "Don't even look my way unless you sell unicorn hair, which is the only thing I really want to buy today". Sounds odd, but imagine if all the salespeople obeyed the sign.

> **"If websites knew your interests, they could avoid showing what you ignore."**

That would be a miracle indeed. Is that really what APML can do?
It's exactly what APML promises: a standard and largely automatic way for you to tell any website you visit what really interests you.

How does APML find that out?
It picks up on all the hints you give during your time online: what you bookmark, your frequently visited websites, the categories in your blog or Flickr account, the topics in a forum you read, what you search for and so on.

What does APML do with that data?
If all the websites you visited knew about your preferences, they could avoid showing you information that you would almost certainly ignore. APML offers a way to collect and encode that information once – in a semi-automatic process that you can control – and provides it in a manner that any website can grab and understand transparently.

Didn't you hint that APML is good for more than advertising?
APML is a portable way to explain to a remote computer what interests you – it can interpret that data in any way it wants. If you think about it like that, you'll see that this can provide you with much more than just personalised advertising. In fact, it could enhance your news browsing, online shopping and research.

Let's start with shopping.
OK. Let's say that Paul loves Vespa scooters so much that he has bookmarks about everything Vespa, is an active member in a forum of Vespa fans and so on. What if Paul then decided to check the website of a used scooters dealer? If that website knew about his Vespa obsession, it could automatically put all relevant listings at the top of its home page.

Cool – and you say the same thing could happen with news too?
Of course. Currently, you choose which RSS feeds to download, but you have to download all the news in each feed, even if you only care about one or two items on a regular basis. Using APML, any RSS front-end or news portal could place the headlines you're likely to want to read first at the top of their listings.

What about research?
If you open up, APML -style, to a search engine or library database, it can factor in your preferences as well as the search string you typed. This helps it give priority to the search results that you'd find particularly relevant.

All right, I'm ready to go technical now. How does APML really work?
In short, by storing your Attention Data into a personal Attention Profile.

I'm still confused – what exactly is Attention Data?
It's a set of variables that contains information on what you like to do online, based on your frequently visited websites, RSS feeds, blog tags or posts, bookmarks, the pictures and/or videos you publish or watch on Flickr and YouTube, the music you listen to at Last.fm and similar portals, and so on.

Then what's an Attention Profile?
It's a structured collection of Attention Data that also gives each section a ranking value – a numeric relevance – and a timestamp. It's the way your interests are logged and ordered in an Attention Profile that makes it possible for websites to understand what interests you right now.

What's the format of an APML file?
XML. When you make use of an APML-enabled service or application, what really happens is that you're either creating or updating an APML XML file, then passing it to a website so it can serve you better.

OK, I think I'm getting it, but what specifically lives inside an APML file?
Your implicit and explicit interests, their sources, authors and the numeric values that rank your interests (as a percentage). Also, the timestamps of each field.

Isn't all this a bit too static? My shopping needs vary a lot over time.
That's what the numeric value and timestamps are for. If your car needs replacing, you'll probably start browsing a lot of car-related websites until you've brought a new one. During that window of time, the value assigned to your interest in cars will continuously increase and its timestamp will be refreshed several times a day. Those high values and fresh timestamps would be a clear indication

for all sorts of services to send car-orientated content your way.

And what will happen after I've bought a new car?
You're likely to stop visiting car dealers online. This will cause the timestamps and values of that interest to decrease or stale enough that any APML-aware service will stop showing you car-related information.

How big is this thing? Do many online businesses already use APML?
The list of websites that already use APML or have declared their interest in it includes Google, Digg, Bloglines, NewsGator, LastFM and Delicious. There are WordPress plugins to create an APML file out of any existing blog on the software (**http://wordpress.org/extend/plugins/apml**) and there are also social portals such as **www.particls.com**, **http://engagd.com** and **www.cluztr.com** that have been built from the ground up to use and share user's Attention Data.

I didn't know any of this. Does that mean APML will become universal, or is it destined to fade into obscurity?
Interesting question – giving an answer is still tricky at this point. Most of the buzz around APML happened around the end of 2007 and the beginning of 2008. After that period, however, things went much quieter. This may either mean that everything is now ready for you to take advantage of APML, or that some part of the plan has gone awry.

What could have gone wrong?
In order for APML to really work as advertised, it must be complete, current and almost invisible to the end user. These are goals in direct contrast with each other and thus will be difficult to achieve simultaneously. Privacy concerns are another issue that may slow down APML adoption.

Right, we haven't mentioned privacy at all yet. Doesn't APML destroy it?
Well, if you're a social network addict or a heavy user of webmail services, such as Google, Hotmail or Yahoo, why worry about reformatting what others already know? Privacy problems could arise from how the data you provide is used by whoever receives it, but that's also true without APML.

Could you be a bit more specific?
Exposing yourself with APML rather than a CV written in Klingon doesn't provide any more control over the receiver, who could resell or misuse that data. Besides, websites that ask for APML files could merrily continue to use other methods to mine more information than you meant to give when you adopted APML.

Is there any other danger?
Yes: APML is another technology to help us see only what we want to see. Ignoring what doesn't interest us is terribly convenient, but may not always be the smartest thing to do. Filtering news, for example, won't change what happens in the world and it may hurt you to be out of the loop. Ultimately, that's a trade-off you'll have to decide about for yourself.

Who develops APML? Can I help?
The APML Workgroup, located at **www.apml.org**. Besides a wiki, the workgroup is also present on Facebook and in a public Google group (**http://groups.google.com/group/apml-public/topics**).

Where can I read more?
There's a FAQ at **http://groups.google.com/group/apml-public/web/apml-faq** and you can find the APML specification at **http://apml.pbworks.com**. If you just want to get a more in-depth idea of what APML can do for you, check out **http://tinyurl.com/3d2h9d**, **www.cleverclogs.org/2007/10/basics-of-atten.html** and **http://eliasbizannes.com/blog/2007/10/explaining-apml-what-it-is-why-you-want-it**. LXF

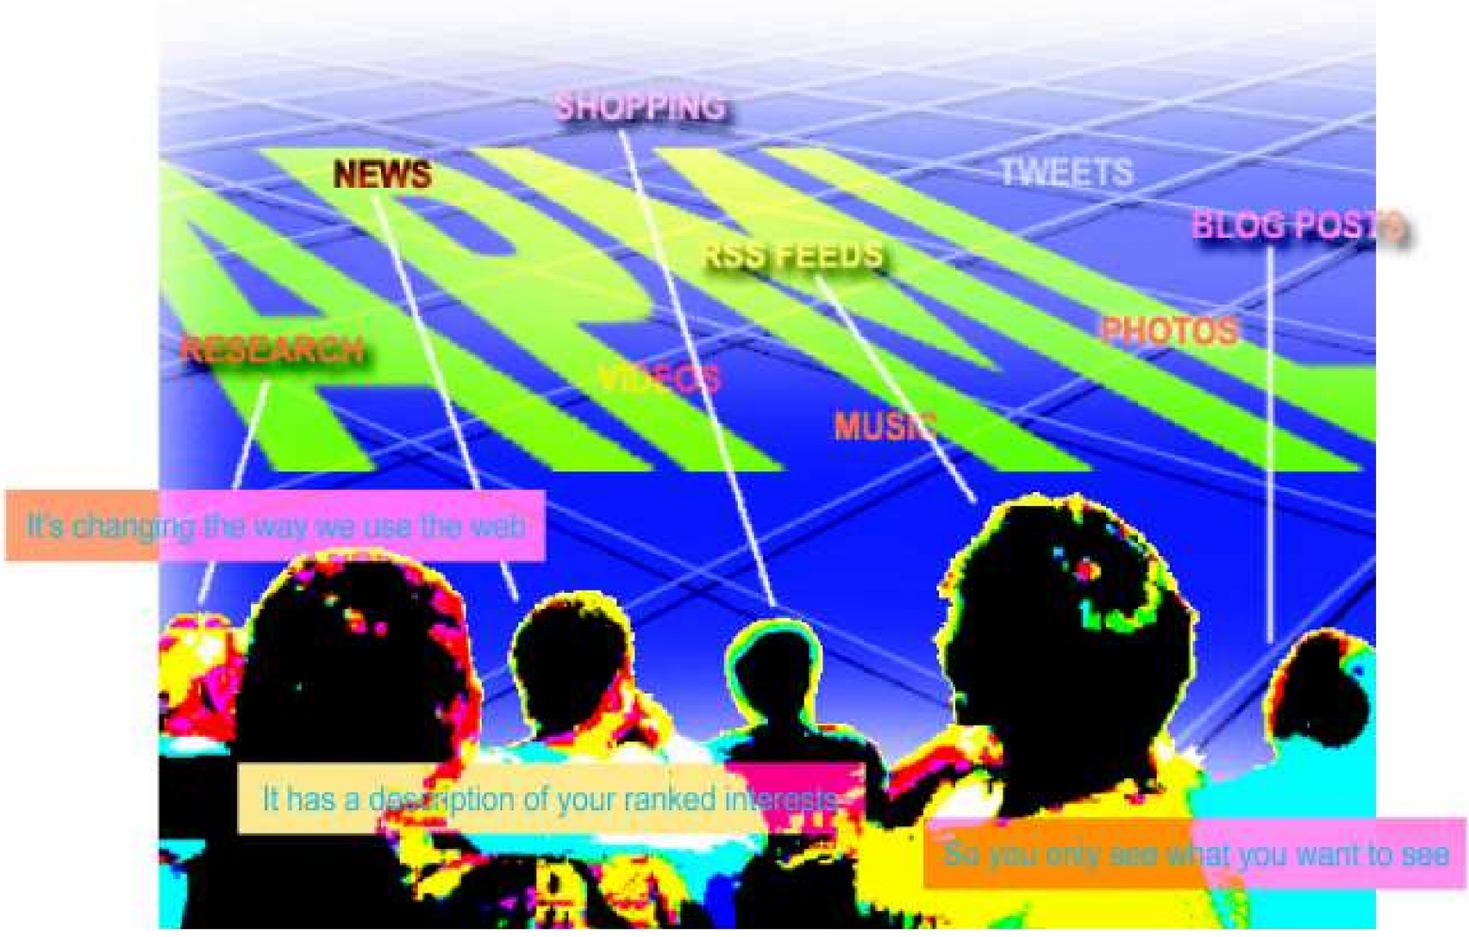

Dr Brown's Administeria

Dr Chris Brown
The Doctor provides Linux training, authoring and consultancy. He finds his PhD in particle physics to be of no help in this work at all.

Esoteric system administration goodness from the impenetrable bowels of the server room.

Take a byte out of global warming

As a reader of this magazine, I'm sure you'll be aware of the need to take care of our wonderful blue planet and of the benefits of recycling. One aspect of our duty to take care of the Earth that has long caused me anxiety is the number of bits that we simply throw away when we're done with them.

So, I'm delighted to announce a new initiative to reduce digital waste. Using the following measures, it's estimated we could achieve a 75% reduction in bit usage by 2015, despite the current rate of growth of the digital market. From 1 April 2012, for example, it will be illegal to create a new TCP, UDP, or IP header. You will be required to recycle old ones and adjust the fields as necessary. To help with this, local authorities will be setting up bit recycling servers in your neighbourhood. Consumers will be required to separate their digital refuse into separate files – Ethernet frames, protocol headers of various sorts, public SSH keys, expired X509 certificates and so on – to send to the recycling servers. Anything that doesn't fall into a designated recycling category must be separated into ones and zeros and the resulting bit streams uploaded for recycling.

If we don't do this, mankind is going to run out of bits and the digital world will come to a screaming halt when the very last one is used. Or maybe it will be the very last zero – it's hard to tell.

The marketing people will no doubt continue to jump on the eco-bandwagon, and I predict we'll start seeing shrink-wrapped software in the stores with a green flash across the corner saying, "This software is made from 90% recycled bits". A glorious change indeed.

This column contains 13,500 bits and represents 15% of your recommended daily allowance.

Easy PC integration

Likewise Open Take the pain out of joining your Linux server to an Active Directory domain.

The world of integration owes a lot to the efforts of the Samba team. However, configuring *Windbind*, *Samba*, and PAM to join a Linux system to an Active Directory (AD) domain still isn't easy. If this is on your to-do list, take a look at the products from Likewise. Its free entry-level offering is called *Likewise Open* and enables you to join your Linux system to Active Directory domains in a single step from the command line or via a GUI. The software performs all the necessary system configuration to enable PAM-aware host services to authenticate AD domain users and honour the user's existing set of domain group memberships.

The early years

Years ago, I attempted to teach a course on Unix and Windows NT Integration. It was ahead of its time – at that point there were no good solutions for resource sharing or single sign-on in mixed Unix, Linux and Windows environments.

Loosen the tie that binds

Once you have successfully joined a Linux machine to an Active Directory domain you can log in using any valid AD user. You will need to enter the username in the format **'domain\username'**. For example:

```
ssh 'example\chris'@hostname
```

Likewise Open doesn't require extensions to the AD schema. The Linux-facing part hooks into the *resolver* libraries and PAM, while the Windows part acts as a Kerberos 5 and LDAP client for authentication and authorisation.

One potential problem with *Winbind* is that it automatically assigns Linux UIDs and GIDs to user accounts held in AD. If you want control over these UIDs and GIDs, the *Likewise UID-GID Module* (which isn't free) extends Microsoft Active Directory so you can assign specific UIDs and GID to Linux users and groups.

See **www.likewise.com** for more details.

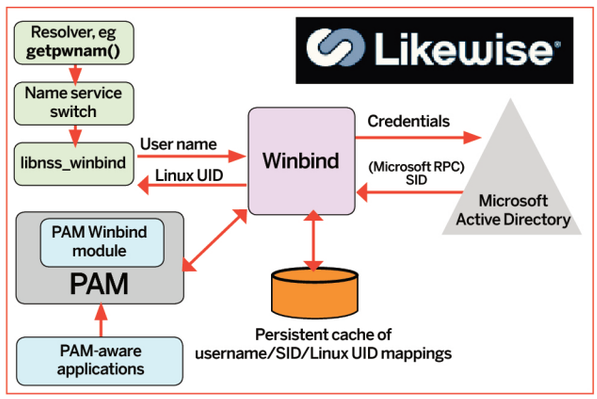

› ***Winbind*** **helps bridge the gap between Linux and Active Directory.**

The really secure shell

SSH Worried by all those failed SSH login attempts that riddle your log files? Batten down the hatches with our six-point plan.

If you run an SSH server on a machine with an internet-facing network connection, it will receive hundreds or thousands of failed login attempts every single day. This may seem phenomenal, but if you don't believe me, look in the log files. But apart from feeling vaguely nervous, what are you supposed to do about them? Here, we'll look at six ways to make your SSH logins more secure.

1 Choose strong passwords

As a first line of defence, strong passwords are vital. Educate yourself (and your end users, if there are any) about what makes a strong password. There's plenty of help online, but in essence you need to use passwords that are long, contain a mix of characters, punctuation and symbols, and avoid words taken straight from the dictionary. If you have a population of users to cater for that are free to set their own passwords, consider using PAM modules such as **pam_cracklib** or **pam_passwdqc** to enforce password strength.

2 Disable direct root login

There is a parameter called **PermitRootLogin** in the *sshd* config file that, if set to **no**, prevents direct root logins on the machine. This makes it harder for an attacker, because he will have to guess both an initial username and password, then crack the root password as well. This defence is raised by default in Ubuntu, although for a different reason – the root password is locked. However, there's a small amount of inconvenience to endure for this on the client side, because you'll have to log in under a regular account first and then switch user to root to use a root login on your server.

3 Disable password-based logins

Although SSH supports several authentication mechanisms, only two are usually enabled by default – password authentication and RSA public/private key authentication. By setting **PasswordAuthentication** to **no** in the config file you'll force *sshd* to use RSA. Before you do that, though, you'll need to copy your public key on to the server. This is easy enough: if you don't already have a public/private key pair, just run **ssh-keygen** to generate one (and set a passphrase) then use **ssh-copy-id** to copy the public key into your **authorized_keys** file on the server. This is a really strong solution, but you'll have to carry your private key around if you routinely log in from many different computers.

4 Run SSH on a non-standard port

You can change the port that *sshd* listens to by adjusting the **Port** parameter in the server's config file. Change it to a 'random' high port number. You will, of course, need to specify this port when you log in, for example:

```
$ ssh -p 22505 chris@myserver.example.com
```

Alternatively, you can configure this on a per-host basis in the SSH config file (**/etc/ssh/ssh_config**).

Security purists would wish me to point out that you can't achieve total safety through obscurity. If you were determined to break in to my machine, you could easily find the port that *sshd* is using with an *nmap* port scan. But most scripted brute-force attacks on SSH will only use port 22. I would guess (and it is just a guess) that moving SSH away from port 22 will reduce the number of break-in attempts by at least 90%. Does this make the machine secure? Not in an absolute sense. Does it reduce the risks? Yes. Nevertheless, this is probably the weakest suggestion of the lot.

5 Ban persistent offenders

I recently came across a program called *fail2ban*, which locks out login attempts from sites that make repeated invalid login attempts. The idea behind the program is fairly simple. It peers into your log files from time to time, (using some fancy regular expression footwork) identifies the lines that report failed login attempts and extracts the IP address of the machine that made them. After a threshold number of failed attempts has been reached, which you can configure, *fail2ban* adjusts the *Iptables* firewall rules to block the offending IP address entirely. This ban can be permanent, or revoked after a specified time.

6 Rate-limit the connections

The idea here is to limit the rate at which any one IP address can establish new SSH connections. If this limit is exceeded, blacklist the IP address (on the assumption that it's an attacker) and subsequently **DROP** all connections from that IP for an hour or so. All of this can be accomplished using the *Iptables* recent module. Note that this solution works entirely at the IP layer and does not distinguish successful SSH logins from failures – it simply limits the rate of attempts. There's a good discussion of this technique at **www.sollers.ca/blog/2008/iptables_recent**.

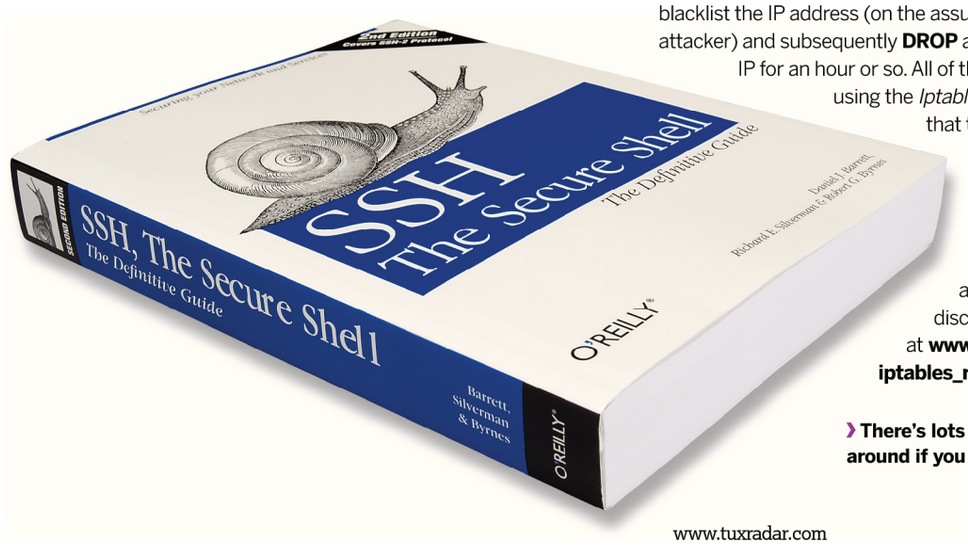

› **There's lots of SSH reading material around if you want to find out more.**

Wireshark

Wireshark Ever wondered what's really happening on your network? Get started with Wireshark and find out...

Wireshark is one of my favourite tools. I think it's because of the power of discovery it confers. I still remember the sense of wonder I had the first time I used it (in those days it was called *Ethereal*) as a whole new world suddenly became visible to me. I felt a bit like Anthony van Leeuwenhoek must have done the first time he peered down his microscope and saw things that no man had ever seen before. OK, maybe I'm getting a bit carried away, but it's certainly both a useful diagnostic tool and a fascinating aid to learning how things work.

Essentially, *Wireshark* is a tool that captures network traffic from one or more network interfaces. It can apply a combination of filtering rules to decide if the packet is of interest and it can pick those packets apart to reveal them in an astonishing level of detail. It can also save packet captures to files and load them back again. Normally, *Wireshark* places the interface into what's called Promiscuous mode, meaning that it will pick up all packets, not just those addressed to its own MAC address. Setting Promiscuous mode requires superuser privileges, so *Wireshark* is usually run as root.

Let's start by taking a look at a simple example of a *Wireshark* trace that's the result of browsing to a web page on a server at 192.168.1.67 from a browser on a client at 192.168.1.69. The top part of the display (*in green in the image at the bottom of this page*) shows the entire packet exchange. Each row is one packet. Packets one to three show the TCP connection handshake, packet four is the **HTTP GET** request and packet six is the reply. Packets 7–10 show the handshake as the connection is dismantled at both ends. The Time column is the time in seconds since the first captured packet and can be useful when diagnosing, for example, sluggish behaviour due to DNS timeouts. In our case, everything is over and done in less than 3ms.

In the screenshot, we've selected packet four for detailed examination. In the middle pane of the display we see a summary of the packet headers for this packet, which correspond to each layer of the protocol stack. The little arrows on the left can be used to expand each layer to show more detail. In the screenshot, we've expanded the application level header – in this case an HTTP packet. We can now see that it's an **HTTP GET** request, and we view the fields of the HTTP request header.

The bottom pane of the screenshot shows the packet's contents, byte by byte, in hexadecimal and ASCII. The section of this display that's highlighted shows the field of the HTTP header that's selected in the panel above – in this case, the **Host:** field.

Capture filters

To do this	Use this filter
Capture only traffic to or from a specific IP	host 192.168.1.44
Capture traffic to or from a specific subnet	net 192.168.1.0/24
Capture only DNS (port 53) traffic	port 53
Capture everything except ARP and DNS traffic	port not 53 and not arp

Display filters

To do this	Use this filter
Show only traffic between machines on the local network	ip.src==192.168.0.0/16 and ip.dst==192.168.0.0/16
Show only traffic from a Dell MAC address	eth.addr[0:3]==00:06:5B
Show HTTP requests where the URI ends in foo	http.request.uri matches "foo$"
Show windows-related traffic	smb || nbns || dcerpc || nbss || dns

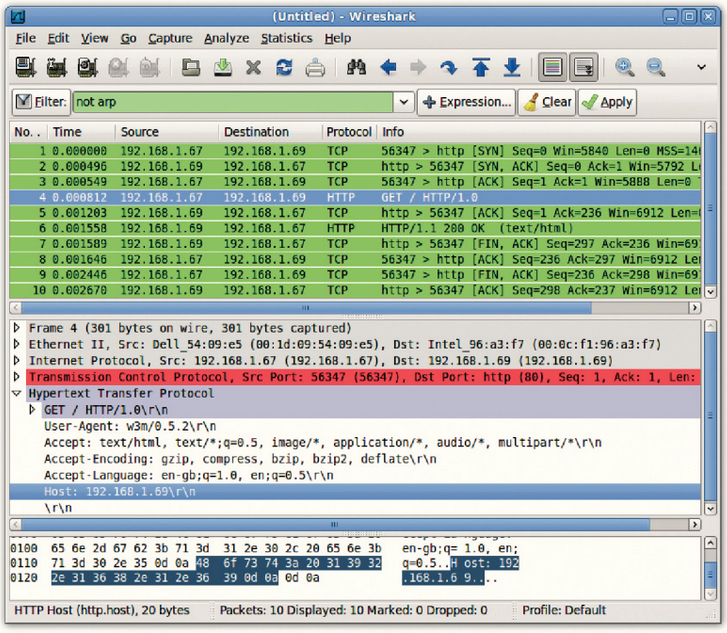

› The main *Wireshark* window. Everything you ever wanted to know about your network in an array of pretty colours.

Filters

Filters are one of *Wireshark*'s most powerful features. A filter is one or a selection of tests on a packet's contents to determine if it is of interest. Filtering is employed at two stages. Capture filters determine the packets that will be retained in the capture buffer, and display filters determine which packets within the capture buffer will actually be displayed. The filter language is rich and you can filter for

Ask permission first!

If you are planning to run *Wireshark* on your corporate network, you might want to ask permission first. Although *Wireshark* is entirely passive and non-invasive (and on most networks it will only show you 'your' packets anyway), corporate policy might prohibit the use of such tools. You have been warned!

Behind the scenes

Behind the scenes, *Wireshark* runs a program called *Dumpcap* to do the work, and *Dumpcap* in turn uses a packet capture and filtering library called *Libpcap*. It is this library that's the real engine powering *Wireshark*. It's also used in other tools that need to capture network traffic such as *Tcpdump*, port scanning tool *Nmap*, intrusion detection tool *Snort* and *Wireshark*'s command-line brother, *Tshark*.

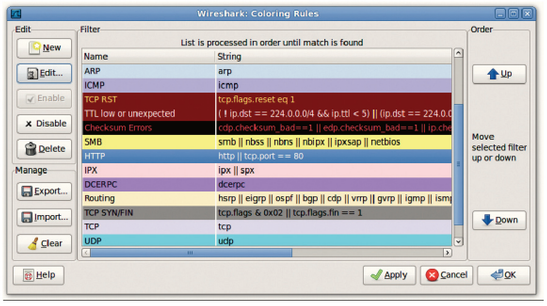

› A colouring rule associates a filter expression with a display colour. For each packet, the first matching rule determines its assigned colour.

practically any field of any protocol, enabling you to home in on the traffic you want to see.

If your eyes are keen, you may have noticed a display filter in use in the packet trace we just examined. It consists of the simple rule **not arp** and I used it to suppress all the gratuitous ARP chatter from my broadband router.

Filters can test for specific protocols such as IP, TCP, UDP, ARP and so on. They can make equality, inequality and numeric comparison tests on the values of fields. For fields that contain string values, they can look for substrings with the **contains** operator and match them against regular expressions with the **matches** operator. They can even match against a specific range of bytes within a field using the syntax **[offset:length]**. Individual tests may be combined using **and**, **not** and **or** operators. All of this makes for a versatile and powerful filtering language.

Wireshark includes a graphical tool to help create display filters. To start this tool, click on Expression in the Filter toolbar of the main screen. For example, the screenshot *below* shows it being used to define a filter that will display only packets that contain an HTTP response code of 404 ("file not found"). The resulting filter rule is **http.response.code == 404**. But there are a great many fields you can filter on, so how do you know what they're called? Well, as you navigate around within a packet in the main *Wireshark* window, the field name of the item you've selected is shown in the status bar. These are the names that can be used in filter expressions. If you look carefully in the status bar of our first screenshot, you'll see the field **http.host** identified, corresponding to the field highlighted in the panel above. There's also a comprehensive list at **www.wireshark.org/docs/dfref**.

The two tables *on page 64* show further examples of capture and display filters, and should give you some idea of the scope of what filters can do. Before leaving the topic of filters, there's a final class of filters to know about, which enable you to specify colouring rules for *Wireshark*. The Colouring Rules editor enables you create new colouring rules (using the same filter syntax), set the colour for that rule and import or export a set of colouring rules. You can also use custom colouring rules to highlight packets while still seeing the background of other traffic. **Wiki.wireshark.org** has a small collection of prebuilt colouring rules to get you started.

Meanwhile, the Capture Options screen enables you to set – no prizes for guessing – the capture options. Here you can select which interface to capture from, set Promiscuous mode, specify a capture filter, and set a limit in terms of packet count, data volume or time.

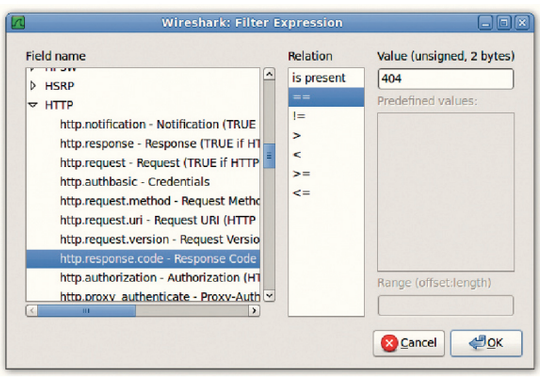

› The Filter Expression editor helps you view the available filter Field Names and build filter expressions.

Doing it on the command line

You can also capture packet traces from the command line using *Tshark*. For example, here we record a packet trace of 100 packets from **interface eth0** to the file **/tmp/tshark.out**, filtering out the ARP traffic:

```
$ sudo tshark -i eth0 -c 100 -w /tmp/tshark.out not arp
```

(The packet filters you can specify here are identical to the ones in *Wireshark*.) Subsequently, we can display the packet trace by reading it back in with:

```
$ tshark -r /tmp/tshark.out
```

The screenshot below shows a fragment of the output from this. The same file can also be loaded into *Wireshark*, for a more graphical view.

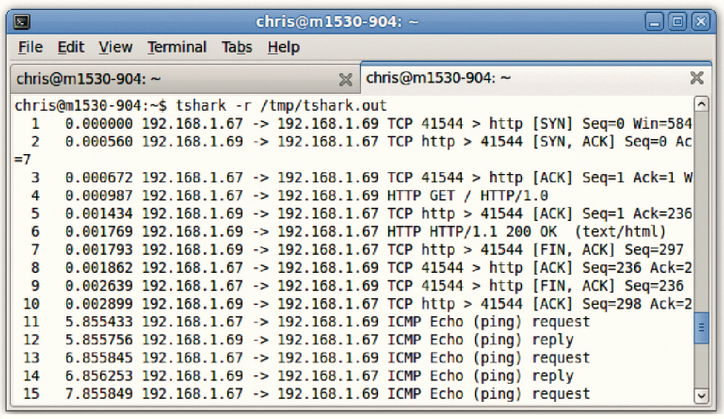

› The command line tool *Tshark* provides the same capture and filtering features you'll find in *Wireshark*, but it isn't nearly so pretty.

Where to learn more

For a formal description of the filter syntax, take a look at the man page for *Tcpdump*. There's also a Wireshark Users' guide, a HOWTO guide, examples of filters and colouring rules and more at **wiki.wireshark.com**, and a through, though excruciatingly detailed filter reference manual at **www.wireshark.org/docs/man-pages/wireshark-filter.html**.

Start here

Three corking operating systems plus apps and games...

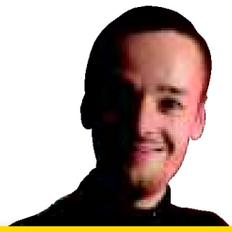

Read on for all the info you'll need to use the **LXFDVD**! If you're new to Linux, open up **index.html** on the disc and go to the Help section for guides to the operating system including:

- What is Linux?
- What is a distribution?
- Booting your PC from the DVD
- Partitioning your hard drive
- Navigating the filesystem
- Root and normal user accounts
- Using the command line
- Installing software
- Getting help online

... and more!

Mike Saunders
New Media Editor
mike.saunders@futurenet.com

Linux distribution

Linux Mint 7

Occupying an impressive third position on the DistroWatch popularity chart, Linux Mint has become a major player in the Linux world. It's built on Ubuntu but has a different interface approach, lots of custom configuration tools, and a development community that's really easy to jump into. We asked **www.tuxradar.com** readers why they loved Mint – here are some of the responses:

"I love this new distribution. It looks better than Ubuntu and has smoother operation. All menus are more intuitive and the distro looks more professional than many of the others." – Carlicus

"I really like how open Linux Mint's development team are to contributions. I added functionality to *MintUpload*, re-packaged, and was granted permission to upload it to the Community repo within three days, even though this was my very first contribution to FOSS. Those improvements have now been accepted into this new release!" – emorrp1

"I have experienced frustrations with Ubuntu which I did not with Mint. I like the Mint menu, and Mint is much more pleasant to look at." – mjjzf

You can try Mint straight from this month's DVD; it runs straight from the disc, but if you want to keep it on your hard drive you can install it by following the guide below. Mint requires:

- **CPU** 1GHz x86
- **RAM** 256MB (beware of shared video RAM cutting this down!)
- **Hard drive** 10GB space

See **user_guide.pdf** in the Distros/Linux Mint section of the DVD for more information on Mint and the extra utilities that make it unique. The version on the **LXFDVD** is 32-bit but will also run happily on 64-bit machines. Also note that this is the Universal Edition which omits codecs and proprietary software that cannot be distributed in every country. See **www.linuxmint.com** for more information.

Behold the new Eco Disc!

If you've already taken the DVD out of the wallet (on the back page), you may have noticed that it feels different. We've moved to a new DVD format called Eco Disc that's better for the environment, using fewer materials and less energy to produce. It's also much more flexible and therefore less prone to warping.

Step by step: Installing Linux Mint 7

1 Boot

Boot your PC from the DVD and hit Enter at the boot menu (see the Help/New to Linux section on the DVD for a guide to changing the boot order if necessary).

2 Desktop

After a few moments you'll arrive at the desktop, where you can try out the software. You can then double-click Install on the desktop to copy Linux Mint to your hard drive.

3 Install

When the installer appears, choose your language, keyboard layout and location. You can always click the Back button to return to the last step if necessary.

How do I...?

When you've installed Ubuntu from the **LXFDVD**, you'll no doubt want to learn more about using the distro. Here's some of the most common tasks you'll want to do and how to achieve them...

» **Browse the web** Click Menu (bottom-left) and then Firefox to start up the most popular open source web browser that's supported by hundreds of great extensions.

» **Chat online** Go to Menu > All Applications > Internet > Pidgin. This program supports every popular protocol, including AIM, ICQ, Yahoo, MSN etc.

» **Edit documents** Click Menu > All Applications > Office and then go choose OpenOffice.org Word Processor (*Writer*), Spreadsheet (*Calc*) or Presentation (*Impress*) – each is also compatible with *Microsoft Office* documents.

» **Edit images** Visit Menu > All Applications > Graphics > Gimp for Linux's most powerful picture editing application.

» **Play music and videos** Head into Menu > All Applications > Sound & Video and try Rhythmbox and Movie Player.

» **Add more software** Click Menu and Software Manager for the super-friendly *MintInstall* tool, or Package Manager for a more advanced program.

» **Open a command line** You'll find this in Menu > All Applications > Accessories > Terminal.

» **Configure the system** Go into Menu > All Applications and then Preferences or System for utilities to change the interface, configure hardware, manage users etc.

Linux Mint's installer is very straightforward, but with the millions of combinations of PC hardware in use, there's always a chance that problems can occur. If you can't boot up to the desktop, try removing any unnecessary hardware (eg scanners, printers) that may be causing trouble with the hardware detection. If you need to get help online, **www.linuxmint.com** has pointers to documentation and mailing lists. You can also try the *Linux Format* forums at **www.linuxformat.com/forums** and Mint's own **www.linuxmint.com/forum**. Good luck!

› Linux Mint's liquorice-and-lime colour scheme makes us want to lick the screen. Slurp.

Don't miss...

MintMenu
This program launcher provides one-click access to your most-used apps and folders.

MintInstall
Via Menu > Software Manager, this awesome utility shows thumbnails of available programs.

4 Partition
Next, the installer will ask you to partition your drive. You can shrink existing Windows/Linux partitions, use the whole drive, or choose manual partitioning.

5 Manual
If you choose to partition manually, create a minimum 10GB root (**/**) partition in ext3 format, and a 512MB swap partition for virtual memory.

6 Create an account
You'll create a user account (your login name and password are case-sensitive); then the Linux Mint files are copied over and you can reboot into your new setup.

More software

DOSes, new apps, games

Linux distribution
Ulteo OVD

If you haven't already seen our review of Ulteo, the brainchild of Mandrake (now Mandriva) founder Gaël Duval, head over to page 28 and see what the fuss is all about. It's been a long time coming, but it certainly looks like the wait has been worthwhile. On the **LXFDVD** in the Distros/Ulteo section you'll find **ovd-iso-latest.iso**, a disc image that you can write to a DVD-R and boot from if you want to install a complete Ulteo system with the Session Manager and Application server. We recommend reading **install.pdf** to get an idea of the process beforehand though, and if you have any problems then try **www.ulteo.com/main/forums**.

Gaël Duval has been beavering away on Ulteo for some time, and the end result is very impressive.

Operating system
OpenSolaris 2009.06

Here's one to expand your Unix knowledge and skills. Solaris is Sun's industrial-strength Unix operating system, typically used on high-end database servers, crunching numbers and also powering a good chunk of the internet. Historically Sun hasn't made any efforts to attack the home desktop market with Solaris, but in recent years the company (now owned by Oracle) has shifted its attention towards open source, and hence OpenSolaris.

> **"OpenSolaris uses many familiar components from the open source world."**

While OpenSolaris has its own kernel, libraries and system tools, it also uses many familiar components from the free software world, such as Gnome and *Firefox*, and therefore isn't a completely alien environment when you first try it. Dive under the hood, though and you'll find a fascinating operating system to explore.

For those who've used OpenSolaris before: this release includes *Time Slider* in the *Nautilus* file manager, which hooks up with the ZFS filesystem and lets you create, delete and browse snapshots of your files, somewhat like *Time Machine* on Mac OS X. There's also *Codeina*, a way to download additional multimedia codecs, along with *Crossbow*, which lets you create virtual network stacks that can be shared on a single NIC. For more information on what's new in OpenSolaris, see the release notes at **www.opensolaris.org**.

On the **LXFDVD**, in the Distros/OpenSolaris section, you'll find an ISO image of the new release that you can burn to a CD-R and boot from. (Note that you have to write it as an ISO image and not just copy the file over – consult your disc burning software's documentation if you're stumped.) If you don't want to burn a CD-R, you can boot the ISO file in *VirtualBox* (see the System section of the DVD).

Sun Solaris is rarely seen outside high-end servers and workstations, but with OpenSolaris it's becoming a solid desktop operating system too.

Important NOTICE!

Before you put the DVD in your drive, please make sure you read, understand and agree to the following:

The *Linux Format* DVDs are thoroughly tested for all known viruses, and are independently certified virus-free before duplication. We recommend that you always run a reliable and up-to-date virus checking program on *any* new software.

While every care is taken in the selection, testing and installation of DVD software, Future Publishing can accept no responsibility for disruption and/or loss to your data or your computer system that may occur while using this disc, the programs or the data on it. You are strongly advised to have up-to-date, verified backups of all your important files.

Please read individual licences for terms of use.

Defective discs

In the unlikely event of your *Linux Format* coverdisc being in any way defective, please visit our support site at **www.linuxformat.com/dvdsupport** for further assistance. If you would prefer to talk to a member of our reader support team, telephone **+44 (0) 1225 822743**.

You need 512MB of RAM, and the OS boots in Live mode so you can try it out without installing. If you like what you see, open **getstart/index.html** in the Distros/OpenSolaris folder on the DVD for a guide to installation and use.

Office suite
OpenOffice.org 3.1

Next up is the first update in the *OpenOffice.org 3* line, and it includes some features that, frankly, should have been in the code tree a long time ago. "*OpenOffice.org* now uses a technique called anti-aliasing", the website proudly boasts, but in reality it just means that *OOo* has caught up with RISC OS of 15 years ago.

But we're just mocking in jest – there are some useful new features in 3.1 as well, and you can read all about it in our review on page 32. To install it, we first recommend removing any existing installations of the suite via your distro's package manager. Search for 'openoffice' and remove all the packages that you see.

Next, pop into the **Desktop/OpenOffice.org** folder on the **LXFDVD** and copy the appropriate **.tar.gz** file for your distro to your home directory. If you're running Fedora, Mandriva, OpenSUSE or another distro that uses RPMs, copy the file with **-rpms.tar.gz** at the end. If you use Debian, Ubuntu, Linux Mint or any other Debian-based distro, copy the **-debs.tar.gz** file instead.

Next, open up a terminal window and extract the archive you copied, eg:

```
tar xfvz openoffice.org-3.1.0-debs.tar.gz
```

Then switch into the resulting directory:

```
cd openoffice.org-3.1.0-debs
```

Next, switch to the root user (**su** or **sudo bash**) and then, if you're using Deb packages, enter:

```
dpkg -i *.deb
```

Alternatively, i f you're running an RPM-based system, enter:

```
rpm -i *.rpm
```

These commands will install the packages. (For RPM systems, also look in the desktop-integration subfolder for menu entries for various systems.) Now you should find *OOo 3.1* in your program menus – if not, you can start it from the command line with:

```
/opt/openoffice.org3/program/soffice
```

Desktop apps
Paperbox and Me TV

More and more we're relying on alternatives to the traditional file manager for organising our documents. We need desktop search engines that actually poke inside files to discover their contents (rather than guessing from the filename), tags and other cool features. *Paperbox* is based on the *Tracker* engine and lets you navigate and organise office documents, text files and other files with thumbnail views and tags. An especially cool feature is the tag cloud, in which tags grow in size as they're used more often. Grab it from the Desktop section.

Me TV, meanwhile, is a digital-television viewing app for DVB devices. It works with DVB-T, DVB-C, DVB-S and ATSC cards, and uses the *Xine* video player for rendering. Alternatively, you can compile it against *VLC*, *MPlayer* or *GStreamer* to fit in perfectly with your favourite choice of media player/engine. Along with the source code, the developers have released an Ubuntu binary package, which is on the DVD.

Virtual machine
VirtualBox 2.2.4

We love *VirtualBox*: it's a great PC emulator that lets you install Linux distros and other operating systems on to a virtual hard drive file without having to repartition your real hard drive. You can use it to boot physical CD/DVD discs or ISO images, so if you want to try Ulteo or OpenSolaris without burning discs, get it installed. It's in the system section of the DVD: there are packages for various distros, but otherwise you can execute the Run file for your architecture as the root user.

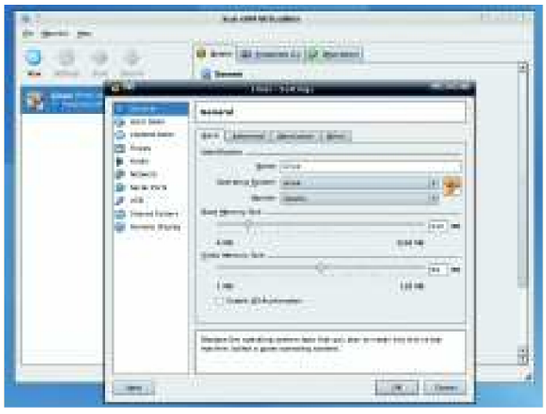

› Among other things, *VirtualBox*'s config window lets you control how much memory the virtual system can use.

Internet tools
BaShare and Twitim

BaShare takes a remarkably straightforward approach to file-sharing. It's a GUI program that opens up a HTTP (web) server on your local machine, then lets you select files so that you can give links to friends on the net. This is really useful if you're having a natter with someone online, and quickly want to send them a file, but your chat software doesn't support it (or the file is too big) and you don't want to mess around with a P2P client. Best of all, you can fine-tune the bandwidth used to make sure that the person downloading doesn't completely hog your connection.

Twitim is a simple Twitter client written in Perl with a *GTK* front-end. Among its features are custom watchlists, sounds, pop-up notifications and XMPP support. We like the clean, non-intrusive design and it's well worth trying out if you're well into the micro-blogging craze.

› *Twitim*: get Twittering with this well-designed micro blogging client, whatever your language!

And finally...
Four excellent games

It's amazing to think that the original version of *Tetris* was written 25 years ago, and yet to this day it's still one of the most popular games on the planet. *Quadra* is a very polished *Tetris* implementation with a stack of tempting features: internet multiplayer; integrated CD music player; 'remote watch' windows (so that you can see what your opponents are doing) and more. Multiplayer *Tetris* is always a riot, so ping your friends online to install it and get playing.

But! If you're bored with regular *Tetris*, try *Tubularix*, which is (as you've probably guessed) "*Tetris* from a tubular perspective". If you're having a hard time imagining this, visit the game's website (**http://tubularix.sf.net**) where you'll see a video of the game in action. At heart it's still very much regular *Tetris* – blocks slide down, and you have to try to make them fit. But because you can move the blocks all around the outside of the well, you have to think differently, a bit like writing with your other hand.

› It's *Tetris*, Jim, but not as we know it.

Also in the Games section we have *Stendhal*, a multiplayer online role-playing game that's built on the *Arianne* game engine and is visually heavily inspired by early 2D *Final Fantasy* games. Then there's *MAX Reloaded*, an open source implementation of the Interplay classic *Mechanized Assault and Exploration* from 1996. On the DVD you'll find packages for Ubuntu and Fedora, or you can compile the source code using the **./configure**, **make** and **make install** (as root) routine. See the Help/New to Linux section of the **LXFDVD** for help with compiling software. LXF

The best new open source software on the planet

LXF HotPicks

Flush » Ubuntu Tweak » Gnome Schedule » Cactus Jukebox » Atomic Worm » SuperTuxKart » Gentoo » BoPlanets » Geeqie

Andy Hudson
When he's not pretending to fix large email networks, Andy dives off coral reefs, checking oyster shells for HotPicks pearls.

BitTorrent client

Flush

Version 0.5 **Web** http://sourceforge.net/projects/flush

We've covered a fair few BitTorrent clients in these pages before, so we can tend to err a little on the harsh side when we see newcomers. The reason for this is that it's all too easy to churn out an average client with little to distinguish it from the crowd. Thankfully, we've just encountered *Flush*, which has temporarily quenched our cynicism.

You'll often start *Flush* by downloading a TORRENT file and opening it in the client. After that, you'll see the download's progress with numerical indicators telling you how big the file is and how much you've downloaded so far. Using the buttons along the bottom of the screen you can dig further to get information on the number of peers that you are connected to, or are connected to you, as well as IP addresses and clients. You can easily find most of this information in a shortened form by clicking the Statistics button, which also enables you to reset the counts if you're particularly fussy about tracking performance. The interface itself is pretty simplistic, and the developer has obviously taken time to consider the end-user experience.

> **"Rounding off the features is the ability to create a Torrent file."**

» *Flush* is a client that can make light work of managing your torrent needs.

Flush with options

There are also a number of options for customising *Flush* – for instance you can specify a working directory and a final destination directory, which is handy if you want to ensure that only finished downloads appear where users can access them. You can also elect to have *Flush* listen to specific port ranges, and there's an area for automation that enables you to configure *Flush* to work on a Torrent file as soon as you copy it to a location. This is good, but it would be even better if you could set specific times when your downloads should occur. Rounding off the feature set is the ability to create your own Torrent file to distribute data. The interface to do this is straightforward – simply point *Flush* at the files to share, add one or more tracker locations and hit Create.

Dependencies-wise, you'll need *libtorrent*, *libboost* and *gtkmm* to get this neat little client going.

Exploring the Flush interface

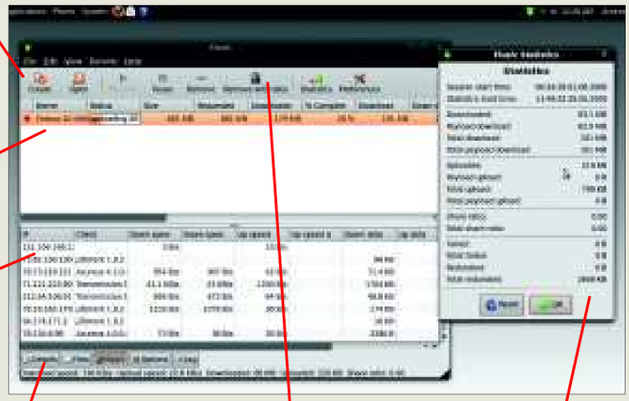

Creation
Click here to start the process of creating your very own torrent.

Progress
Headline figures about current downloads are displayed here.

Peers
Like most clients, *Flush* details who's currently working with you.

Icons
You'll find more information about your torrenting activities here.

Delete
Downloaded the wrong files? This will delete the torrent and its data.

Statistics
Get in-depth information about how your download is progressing.

Configuration tool

Ubuntu Tweak

Version 0.4.7 **Web** http://ubuntu-tweak.com

One of the classic key strengths of Ubuntu has been its relative ease of use – and while other distros have played catch-up, Ubuntu has matured into the most popular flavour of Linux currently available. Due to this position of strength, an even greater emphasis has been put on providing tools and utilities to assist users in getting their installation just right, and it's in this spirit that we came across the appropriately named *Ubuntu Tweak*.

It's a package designed to give you direct access to some of the settings that you would normally have to hunt around for, much as *TweakUI* used to do for Windows. Covering a wide range of settings – including controlling startup behaviour, desktop and personal options – *Ubuntu Tweak* is a great way to customise Ubuntu to fit your tastes exactly without spending hours researching on Google. It even gives you access to some of the more popular alternative repositories, such as Google's Linux Repository and the Opera repository.

Pimp my distro

Once you've installed the program, you can find it in the System Tools menu under Applications. Open it and you're greeted with a short intro screen giving you highlights of the program's functionality. In the main window, each group on left-hand side expands to give you options. For example, under Applications you'll find Add/Remove, Source Editor, Third Party Software and the rather useful Package Cleaner, which helps you reclaim disk space taken up by redundant packages and caches – something that you'd normally need the command line for. *Ubuntu Tweak* exists within userspace, but you can also make system-wide changes using *policykit*.

> "The useful Package Cleaner helps you reclaim disk space."

Available in a Deb for Ubuntu, or in tarballs, *Ubuntu Tweak* isn't going to appeal to everyone – Fedora users, for instance – but it should find more than a few fans among the new(ish) converts to Linux who have just made the switch to Ubuntu.

Want to eke out that last little bit of functionality from Ubuntu? *Ubuntu Tweak* might be the tool for you.

Music manager

Cactus Jukebox

Version 0.4.1 **Web** http://cactus.hey-you-freaks.de

Our personal music libraries are constantly growing, which means it's becoming increasingly important that we have decent software to manage them. *Cactus Jukebox* is another contender for that title, although its interface seems confusingly cluttered initially.

That said, you do start with a blank library – *Cactus* doesn't force you to have all your music files in the Music directory found in most Gnome-based distros these days. A quick look under the File menu produces the Add Directory entry, enabling you to browse to the top-level folder that holds your music and recursively scanning each folder underneath it to retrieve a listing of all compatible music files – OGG compatibility is standard. On the left-hand side is a list of all the registered artists with related albums nested underneath. You can use the directory tree to browse this, while the tracks themselves appear in the panel above the controls. *Cactus* also includes a quick and handy search function that uses fuzzy logic, helping you find tracks based on track, artist or album name.

Tabbed management

On the far left-hand side, you'll see three different tabs: the first is for your music library, the second for network-related activity such as streaming internet radio, and the third enables you to manage music on a media player by dragging and dropping tracks to it.

If you're planning on using *Cactus* to rip your CDs then you'll need the *cdda2wav* package to assist with the ripping; on top of this you should also ensure that you've got an MP3 encoder installed, along with *MPlayer*.

> "Cactus includes a quick and handy search that uses fuzzy logic."

Cactus is a useful project on a different path to the one trodden by *RhythmBox* and *Banshee*, particularly because it relies upon a sole developer. We think it merits further interest from the wider Linux community, and would encourage you to take a look and lend a hand; we're sure it would benefit from the additional assistance.

Your one-stop-shop for managing your music collection, Cactus Jukebox will handle things from start to finish.

System task scheduler

Gnome Schedule

Version 2.10 **Web** http://gnome-schedule.sourceforge.net

If there's one thing that all system administrators have in common, it's their drive to automate run of the mill tasks so they can focus on more important things. Most sysadmins will be familiar with crontab, the home of the instructions that *Cron* uses to carry out tasks according to a fixed schedule, but at home you're more likely to want to schedule a regular backup, or execute a script periodically. For this kind of scenario, it's unnecessary to delve deep into the command line, because you can use a handy utility that not only gives you power over *Cron*, but also gives you limited access to the **at** command – *Gnome Schedule*.

Scheduled tasks

To get started you'll need to ensure you have Python GTK and Python Gnome bindings and then compile using the trinity of **./configure**, **&& make** and **&& make install** (as root). You'll then find *Gnome Schedule* nestled in among the other entries in the System Tools folder of the Gnome menu. Firing it up presents you with a rather plain screen, but fear not – this is merely your blank canvas to work upon. Now click the New button to be given three different types of task. At the moment, there will be no templates defined (you have to create them, more on that later), so you'll have to choose either a recurring or one-off task. If you elect to create a recurring task, you'll then see a simple dialog box prompting you for a description for the task, along with the command (or shell script) to be executed and a drop-down box enabling you to control how the command handles any outputs. If you're planning to schedule an *X*-based application, you can also set up *Gnome Schedule* to suppress its visual activity, forcing it to work in the background.

Further down, you can set basic timing options, including the option to run the job at reboot (handy for a sweep of sensitive data, for example). Alternatively, you can provide a more detailed schedule using the Advanced area. We especially liked the granularity that the Edit buttons provided. The buttons are found at the side of each of the Advanced fields and with them you're able to be much more precise when setting up the frequency of your tasks. For instance, you can specify that the job should only run during certain minutes, hours or days. At the bottom of the dialog, *Gnome Schedule* handily uses natural language to tell you the schedule the job will work to and you're also able to either commit the job or add it as a template.

> **"We especially liked the granularity that the Edit buttons provided."**

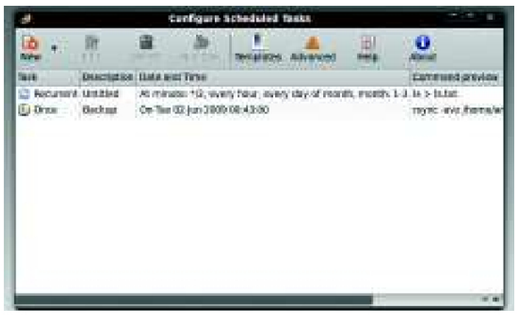

It may not be a looker, but *Gnome Schedule* is comprehensive and easy to understand.

One-off jobs

On the flip side, you can also create a task that will happen at a specific time on a specific date. This dialog box is much simpler, because it only requires you to specify the task to be executed and the date and time at which it should be launched. There's a calendar picker, enabling you to browse through the weeks and months quickly. Again, you're able to save the job as a template that you could use frequently – *Gnome Schedule* allows for the creation of a serious template library, so if you find a specific job works well for you, or even a set of date/time criteria, you'll no doubt make use of the template facility

Ultimately, we found this to be a handy application for home users, but it's unlikely to suit more experienced system administrators.

Step by step: Scheduling a job

Start the process

Clicking the New button presents you with your first choice: is this a recurring job, or a task that's only going to happen once?

Getting it right

Here's the important bit – getting the command correct. Now you can carefully plan the schedule for your job.

The shoulders of giants

If you get into the habit of making templates from existing jobs, you'll have a library to work from in the future.

HotGames Entertainment apps

Arcade game

Atomic Worm

Version n/a **Web** www.charliedoggames.com

› **Your worm is constantly hungry – even when gobbling isotopes. And what's all the floating poetry really about?**

Probably the first game we can remember playing on a computer was *Snake*. If you weren't around in the eighties, or have never owned an early Nokia mobile phone, it's based on a simple concept – you guide your snake around a screen gobbling up treats as you go. The more you eat, the longer your snake gets, but you need to avoid running into yourself or it's game over.

The original could hold your attention for a few hours because it was quite addictive. The problem was, it was also fairly simple; once you'd mastered it, there was no real challenge to keep you coming back.

Atomic Worm takes the venerable game, gives it a darn good 21st century seeing to and boy what a difference 20 years makes. Instead of giving you free rein to take your worm wherever you choose, you're instead guided around a grid where you use your mouse to select your worm's next move. You can only go one move ahead, so you can't line up a series of moves to help you out of a jam.

Fortunately, your worm always moves at a steady pace around the course, so you just have to keep going until an exit opens up. The shape of the grids alternates between levels – you'll play on hexagons, triangles, squares and more – and it's not always fixed.

But it's not just the navigation of your serpent that's changed – you'll also find the developer decided to introduce mines. However, as long as you collect three identical items in a row, you'll get a short speed bump and shield to help you burst through the mines and you'll also shrink a little. These tweaks mean it's a vicious little game that warrants plenty of time to get to grips with, but if you suffer from motion sickness the moving levels will probably make you heave. On the plus side, this will help force you to take breaks.

> **"It's a vicious little game that warrants plenty of time."**

If you want to wrap your coils around *Atomic Worm*, the **install.sh** file takes care of most of the hard work for you, so there's no need to worry about dependencies.

Racing sim

SuperTuxKart

LINUX FORMAT HotPicks Revisited

Version 0.6.1 **Web** http://tinyurl.com/nn4yqv

› ***SuperTuxKart* is a great game packed with cute characters and devilish drives.**

It would be unfair of us to overlook the release of an upgrade to a legendary Linux game. *SuperTuxKart* has long held a special place in our hearts, fuelled by nostalgic memories of the game it's based upon – *Super Mario Kart*.

The premise is a simple one: you scoot along a 3D track trying to beat your opponents to the finish line. You do this with a mixture of skillful driving and devious use of the packages you'll collect along your way. In these packages is a selection of offensive weapons and defensive objects to help you get ahead. It's hard to pick a single favourite: our top two are the large bowling ball that ploughs through opponents and the sink plunger that slows down the opponent ahead of you. The tracks are also littered with treacherous banana peels and speed-inducing nitros.

One of the nice touches in *STK* is the use of the avatars from various open source projects as racing drivers, including a dinosaur based on the Mozilla logo and Tux himself. However, this new version isn't a significant upgrade – a new kart and map make an appearance, but the real change is that *STK* now supports the use of add-on packages to enhance and extend the game further.

> **"SuperTuxKart is an enjoyable interpretation of a gaming classic."**

To get the game running at a decent frame rate you're going to need a dedicated graphics card with the latest drivers. Other than that, you'll need *SDL* and *plib* installed before you compile *STK* from source.

All in all, *STK*'s a really enjoyable interpretation of a gaming classic and a game that will keep on growing.

File manager

Gentoo

Version 0.15.4 **Web** http://obsession.se/gentoo

Lightning fast and highly-configurable: we're not talking about Gentoo Linux here, but an impressive file manager.

Before Gentoo Linux there was *Gentoo* the file manager. Based on the age-old paradigm of a twin-pane file manager, *Gentoo* hits our nostalgia button for many reasons, but primarily because it replicates a popular Amiga-based file manager called *Directory Opus* right, down to the way that it enables you to customise the array of shortcut buttons. You have access to a wide range of useful commands by default, but should you require any additional shortcuts you can add or remove buttons as you wish. You can even colour-code them to suit your style of working. Through the interface you can also quickly add files to compressed archives, as well as browse and extract files. You can also add shortcuts to common locations within your filesystem to make your navigation smoother.

The quick access theme prevails throughout the application, even to the brief summary in the title bar telling you the number of directories and files present in your active pane. Each file type can be colour-coded to help you distinguish them, and you can add actions to specific types as well. You can also use regular expressions to identify the different file types, even going so far as to identify a filename as being in a specific format.

Blazingly fast

Thankfully, all this customisation doesn't slow *Gentoo* down. In fact, the beauty is that it's lightweight and you don't need to satisfy any major dependencies, since it's coded entirely in C. Just fire off a standard **./configure**, **&& make** and **&& make install**. Once you've got into *Gentoo*, you'll no doubt find yourself moving through files and directories a lot faster than you would do if you were using *Dolphin* or *Nautilus*. It has all the speed of a terminal, but with the visual advantage of a GUI interface. You'll need to manually add *Gentoo* to any menu structure you have, but once you've started using it you may find it difficult to go back.

> **"You can use regular expressions to identify the different file types."**

Astronomy tool

BoPlanets

Version 1.3 **Web** http://tinyurl.com/mj8glm

View the skies of the distant future from cities or planets with this intriguing planetary tracking application.

Here at LXF Towers, it's not just free software we admire – we also enjoy the majesty of the heavens themselves from our lair. So it was with interest that we stumbled upon *BoPlanets* – a neat little tool designed to map out the locations of the planets within our solar system.

First of all, you need to know it's a Java application, so it will run happily on any supported platform. On startup, you're asked to select a planet to observe from and the heavenly bodies you wish to track (and no, Sarah Michelle Gellar isn't an option). Clicking Start takes you into the main interface, where, by default, *BoPlanets* displays the precise locations of your chosen planets at the current date and time. That might seem a little boring, but click and hold on the central area and you'll see the planets start to shift in relation to your observation point. By default, the date is stepped forward two days at a time, but the rate can be controlled by the Stepsize field in the toolbar. For our own benefit, we tended to switch off the radius vectors, preferring to see the path of the planets without the inherent noise from those lines.

Star gazing

You can view the orbits of the planets from three different points, or elect to view them as they would appear from ground-level at a city of your choice from around the world. You can also choose to view the planet's paths in a standard elliptical view, or use a logarithmic view to show the relative distance between each celestial body. What's staggering is the sheer amount of plotting this program can do: you can take any step forward from a sliver of a day to more than a whole Earth year to see the paths that the planets will take. And as you do, it's mind boggling to see future skies – learning that the next time Neptune and Pluto will intertwine is around the middle of the 23rd century is fantastic.

> **"What's staggering is the sheer amount of plotting this can do."**

All in all, *BoPlanets* an intriguing concept. Spend a few moments with it and it's easy to lose track of time.

Image viewer

Geeqie

Version 1.0 beta 1 **Web** http://geeqie.sourceforge.net

These days, it's easy to take being able to double-click on an image file and see a either a thumbnail preview or the image itself for granted. In Gnome, for example, you'll tend to have either *Gthumb* or *F-Spot* for this task, and you'll find an assortment of other small and light image viewers across the other desktop environments. But you don't care how it happens – you just want to see the image without having to faff around.

Geeqie certainly falls into the 'not-faffing-around' category for us, from its low dependency needs to the small file browser that enables you to navigate through a large image library easily. The preview itself gets the lion's share of the screen space, along with a selection of information lumped below the image area. Exif information isn't provided by default, but this is easily found under the View menu, providing you with a rather large screen detailing everything you could possibly want to know about a photograph. In fact, *Geeqie* will give you a lot of information, should you want to see it, and you can also add information back by tagging each photograph with keywords or categories. It's a handy little application, particularly for photographers, and a worthy alternative to some of the mainstream options currently out there.

› *Geeqie* is a simple, lightweight image viewer that doesn't get in your way.

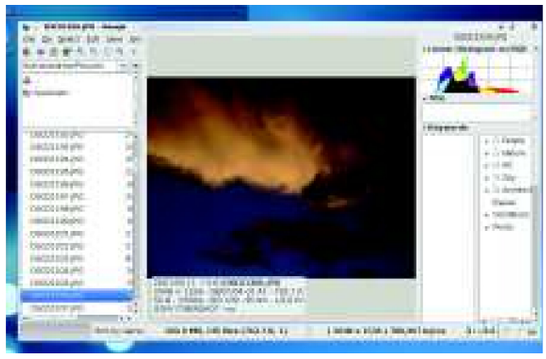

Text Editor

Minimum Profit

Version 5.1.2 **Web** http://triptico.com/software/mp.html

Sometimes it can seem like programs go out of their way to provide you with numerous toolbars and menu options. This can be overwhelming, so we usually prefer simplicity in our working environments.

That's why *Minimum Profit* (*MP*) comes as a relief – it's a text editor designed for use by programmers and developers, and there's not a toolbar in sight. Instead everything is located under five menu entries, leaving the large central area blissfully free for coding, coding and yet more coding. The majority of menu options have keyboard shortcuts, so you may even find yourself ignoring the minimalist menu interface once you've got used to how *MP* works. Code highlighting is based upon the file extension of the current file, which helps keep your work organised, and each file can also be opened in a separate tab.

Minimum Profit requires little in the way of dependencies and will feel at home under any of the major window managers. What's more, despite the simple interface, there's extensive documentation available should you wish to delve deeper. Given that, we'd be happy to recommend you replace *Vi* or *Emacs* with *MP* and, in a few months, we doubt you'll look back. LXF

› *Minimum Profit* gives you a clean and uncluttered work space to code in.

Also released

New and updated software that also deserves a look...

K3b 1.66.0 alpha 2
K3b is CD/DVD writing application that's popular with KDE fashionistas. It has good CD and DVD burning support and a handy user interface, but it's not as simple as *Brasero*.
http://k3b.plainblack.com

› A number of KDE fanatics use *K3b* for their disc burning needs.

Diffuse 0.3.3
This is a graphical **diff** tool with links into all the prominent source code management systems.
http://diffuse.sourceforge.net

Canorus 0.7
Canorus is a music notation editor in a similar vein to *Frescobaldi*. It offers a range of import and export filters, and can make use of *Lilypond* to produce stunning scores.
http://canorus.berlios.de

› Version 4.3.4 of *Eric* has new navigation options and features.

Eric 4.3.4
This venerable IDE just received another patch update, adding a host of new options to an already bursting list of features.
http://eric-ide.python-projects.org

Jacl 2.4.8
Harken back to the days of text adventures with Jacl, a language designed to help you create your very own fantasy adventure.
http://freshmeat.net/projects/jacl

OBM 2.2.3
Looking at alternatives to *Microsoft Exchange*? If so, you should consider *OBM*, a full-featured groupware application with a rich web interface.
http://obm.org

American subscriptions

THAT'S JUST $10.77 PER ISSUE!

Exclusive US & Canada offer!

Subscribe and save 40%

PLUS Get a set of fantastic Linux Format stickers!

BENEFITS OF SUBSCRIBING

- Access over 1,000 magazine PDFs in the exclusive online subscriber area
- Have every issue delivered direct to your home or workplace
- Never miss another issue
- Save $$$

Pay by continuous credit card for **$32.99** a quarter or pay **$131.99** for the year!

Stop paying $16.99 per issue!

LAST CHANCE TO GET THESE EXCLUSIVE STICKERS!

Subscribe online: www.imsnews.com/linuxformat

Or phone TOLL-FREE on 1-800 428 3003. Please quote a006

OFFER ENDS 30 AUGUST 2009

To take advantage of this offer, visit the above URL and enter code a006. Minimum subscription terms is 12 months. You will receive 13 issues of Linux Format magazine a year. Your subscription will start from the next available issue. In the event of stocks becoming exhausted, Future reserves the right to replace with item of similar value. If you have any problems in the next 60 days please notify us in writing and we will refund all unmailed issues. You will be charged in USD.

LINUX FORMAT Tutorials

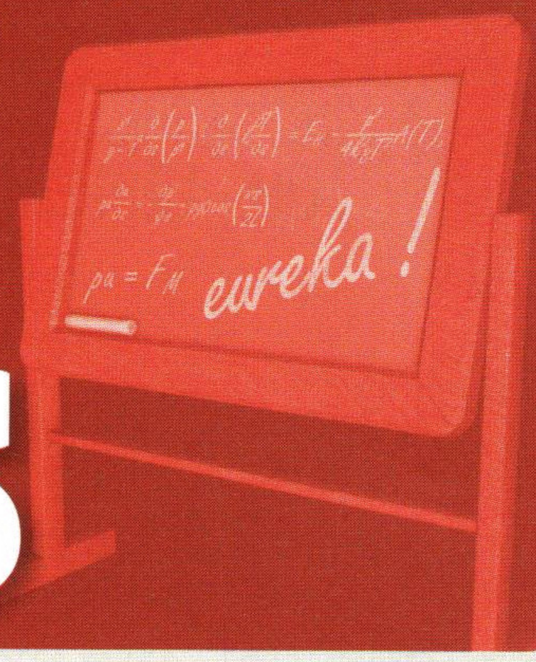

Turn on, log in and let our experts help make you a smarter Linux user

MIKE SAUNDERS has to have a least a hint of Unix in anything he touches.

Pre-dictions

As I write this, technology pundits around the internet are agog to get their hands on the new Palm Pre mobile phone. Not since Apple launched its now iconic iPhone have we seen such hysteria around a device. Yet while everyone coos over the design and flashy interface, I'm getting a warm, fuzzy feeling knowing that Linux is busy working away underneath the shiny exterior. WebOS, Palms's penguin-powered mobile operating system, is doing all the grunt work.

Pocketable penguin

Now, I'm notoriously bad with my predictions, but I do see a very healthy future ahead for Linux on mobile devices. We already have the well-received Android, Moblin is coming along in leaps and bounds, and if the Pre takes off as well as Palm hopes, we could see another batch of Linux-running devices in everyone's pocket. But the big question is: how many people will know that they're actually running Linux? And should they even care? Answers on a postcard please.

mike.saunders@futurenet.com

This month learn how to...

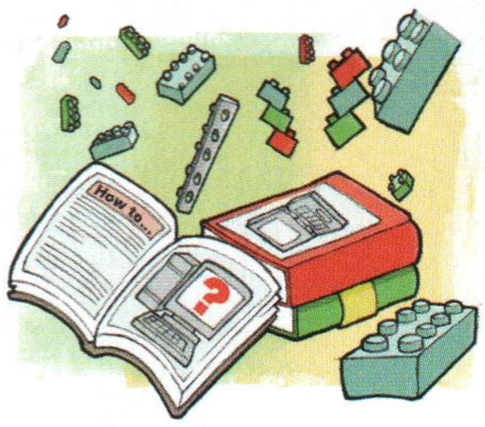

Stream your music ... 78

Don't waste time copying your MP3s and OGGs to every machine you own: **Andy Channelle** explains how to use *Sockso* to share your music around.

Have fun with fire 82

Not in the pyromania sense – but by following **Michael J Hammel's** *Gimp* guide you can add beautiful flame effects to your work with ease.

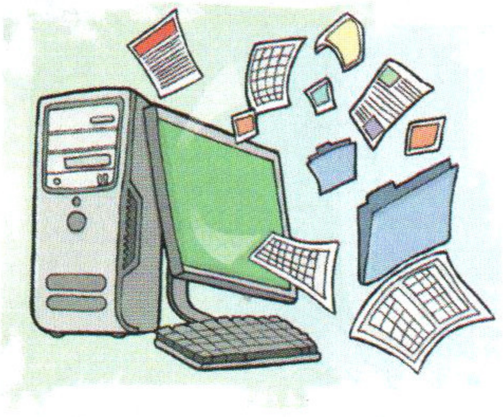

Script your office 86

Generate a stack of invoices with just a click of your mouse. **Marco Fioretti** makes *OpenOffice.org* work for you, not the other way round. Hurrah!

Defeat spam forever . 90

Is junk mail crippling your mail server? Going mad from R0L3X W4TCH messages? **Neil Bothwick** combines *Postfix* with *Dspam* for top results.

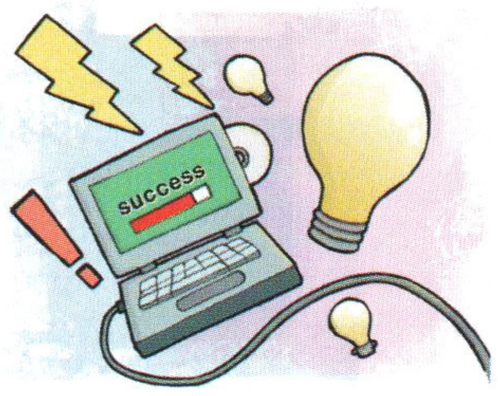

Write a chat bot 94

Python guru **Nick Veitch** has no reason to talk to mere humans – his newly coded Jabber chatbot provides the best conversations on Earth.

Secure your servers.. 98

Stop crackers at the first hurdle: block open ports and set up an intrusion detection system. **Martin Meredith** will show you the way.

Tip of the month: Woof

File sharing with *Samba* or NFS is easy once you've got it set up on both computers, but what if you just want to transfer a file to another computer on the network without the hassle of setting up software? If the file is small, you may be able to email it. If the computers are in the same room, and USB devices are permitted on both, you can use a USB flash drive, but there is also another option.

Woof is a Python script that will run on any Linux (or similar) computer. The name is an acronym for Web Offer One File, which sums it up fairly well, as it is a one-hit web server. There's nothing to install; just download the script from the homepage at **www.home.unix-ag.org/simon/woof.html** and make it executable, then share the file by entering

```
./woof /path/to/myfile
```

It will respond with a URL that can be typed into a web browser on another computer on the network – no software beyond a browser is needed. *Woof* will serve the file to that computer and then exit (you can use the **-c** option to have it served more times). *Woof* can also serve a directory, like so:

```
./woof -z /some/dir
```

This will send a *gzip*ped tarball of the directory, and you can replace **-z** with **-j** or **-u** to get a *bzip*ped or uncompressed tarball. If others like *Woof* and want to use it, you can even let them have a copy with

```
./woof -s
```

Sockso: Build a

Tired of having to transfer music around? **Andy Channelle** shows you how to create a personal music server so you can listen to your tunes anywhere.

Our expert

Andy Channelle
Andy seems to have been taking his first steps in free software forever, and has been interested in technology since the advent of the Dragon 32.

We love music and we love computers. So put the two together and you have the perfect combination. No, not Kraftwerk, but a server that can stream your MP3 collection to any computer you use. In this tutorial, we're going to download and install an open source media server package, configure it to play our library, design some playlists and remotely access these over a local network. We'll also look at the potential for accessing your collection from outside your home network.

The software we're using for this project is called *Sockso* and the latest version, 1.1.8, is available from **http://sockso.pu-gh.com**. *Sockso* is available for Linux, Mac and Windows, and there's no installation as such. Instead all you need to do is unzip the file to a suitable location and double-click either the **sockso.jar** file or the **linux.sh** script. The JAR file extension should warn you that *Sockso* is *Java*-based software, so if you don't have it already, find and install version 1.6 of *Java Runtime* through your package manager.

› **You can access your music via this smart-looking website.**

› **Create your server accessible playlists by dragging songs from the file tree to the Playlist window.**

When you fire up *Sockso* for the first time you'll be presented with the main interface. This has a series of tabs running across the top, a note of your current IP address at the bottom and a two-pane section in between. These will eventually be used to access your library and to construct playlists, but at the moment they both look a little desolate. We need to add some music.

Hey Mr DJ, put a record on

To begin adding tracks, select the Collection tab from the top of the window. Here you'll see a large window and a pair of buttons that are used to add and remove folders from your collection. Click the Add button to launch a file browser and navigate to your music collection. This tool imports subfolders, so there's no need to import albums individually – just choose the parent music folder and everything beneath it will be added to the server. Well, links to tracks are created within the software's database – the files aren't actually copied to the new location.

We added the folder **~/home/user/Music** (where **/user** is the server admin's username) and it took about three minutes to add 3,600 songs to the database. You may find it takes longer to parse larger collections, so you could consider leaping beyond these first steps and installing *Sockso* with a *MySQL* back-end to make it easier to handle big collections.

By default, *Sockso* will scan any folders added every five minutes for new tracks. This seems rather excessive, so you

» **Last month** We got organised with Basket Note Pads and OOo outlines.

music server

can change it to scan every hour by going into the General tab and setting Scan Interval to 60. If you add tracks later and want to update quickly, go to Collection and select Scan Now. Oh, and while you're in the General tab, why not rename the server with a tasteful name, such as The Unstoppable LXF Music Machine?

It's time to test the server, so fire up your browser and type **http://localhost:4444** into the address bar. If everything's working correctly, you'll see the standard *Sockso* interface, which presents you with various links under the splashy header and an alphabetical strip on the right-hand side of the window. Selecting a letter takes you to a list of artist names that begin with that character – to the right of each artist name is the number of albums available and to the left are a trio of icons. The first will play everything under that artist, the second adds the tracks to a playlist and the third will download the track (you won't need to do this, as you're running a local version). Clicking on the artist's name will present you with a list of albums and selecting an album displays the individual tracks.

You can also use the Search bar to look for a particular artist or song. This is dynamic, so the list of available tracks will be pared down as you type.

Making playlists

Now, wouldn't it be great to have playlists that are accessible anywhere in the house? Well, *Sockso* makes it simple. Go back to the Music tab and you'll see your music displayed on the left-hand side. Click the disclosure icon (+) next to each artist to reveal albums and then the icon next to an album's name to see its tracks. From here, you can drag tracks to the window on the right. Songs can be re-ordered by dragging them up or down and deleted with the Delete button.

Once you're satisfied, click Save Playlist and give it a name. To see your playlists, click the Playlist tab above the file tree. To edit one, click and drag it to the right-hand pane. Once you've altered your playlist, click Save Playlist again and type in the amended playlist's name.

Return to the server and click the Playlist link. The playlists you've created will be listed under the Site Playlists section

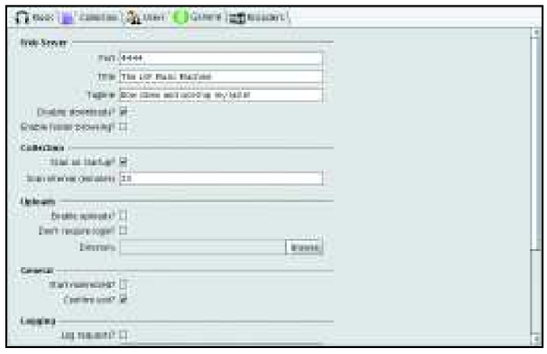

› You can disable downloads to ensure you won't end up with lots of copies of your music everywhere.

and to the right of this is User Playlists, which enables other users create their own playlists via the web front-end.

Their efforts will appear in the Administration GUI within the Playlists tab, so if a remote user creates a particularly great playlist, you can copy it and make amendments. Just drag it to the Playlist window (on the right), make changes and then save it with a new name.

Network access

Now imagine we're in a house where there are four users who want to access the server. If we're being completely open, each person could direct their browser at the local IP address for the server, but this is likely to be different to the address *Sockso* shows. You can find a machine's IP address in a terminal by typing **ifconfig**. Enter this address into your browser, followed by **:4444** – the port the server's running on.

We want to ensure each of our users logs in to access music and restrict downloads. To do this, go into the Users tab and click the Create User button. Provide a username, password (twice) and an email address, then click on Create User. Before leaving this section, click the Require Log In button. Complete the process by going back to the General tab and selecting Disable Downloads. You've now created your own home network radio station. Smashing!

Quick tip

If you're up for a bit of code hacking, you can find all the design information for *Sockso* in the **application/resources** folder.

› The pop-up player is functional without being too bland.

Playback

By default, any playlists and files you select to play will be opened in a pop-up *Flash* player. This worked on every machine we attempted to connect on, regardless of operating system, and also enables you to start a long playlist and then navigate away from the *Sockso* web interface without losing your music. However, if you have a problem with Flash, there are other options including streaming to various applications, such as *iTunes* and *Windows Media Player*, or using an embedded player.

The *Flash* player displays the playlist, enabling you to select tracks to play, a set of transport controls for navigation and any associated album art for the track currently playing.

Basket: Order

Basket can do more than to-do lists – it's also an ideal research tool.

Last month, we built a basic *Basket* system to deal with to-do lists and simple notes. Now we want to revisit this powerful application to look at other ways it can be used to organise, manage and share information. Specifically, we're going to look at the software as a research tool, which could be useful for writing a novel, preparing a report, designing a presentation or planning a household project. We'll also look at acquiring, editing and organising notes, importing information from a variety of sources, sharing information with friends or colleagues and using *Basket Note Pads* to present your research to others.

Quick tip

Save your baskets file in a *Dropbox* (**www.getdropbox.com**) folder so you can access them from any computer.

Preparing the app

First, create a new basket, which can be done by right-clicking the Baskets list and selecting New Basket, or by pressing Ctrl+N. Customise it using the various options – we've gone for a freeform basket – then provide a suitable name and click OK. The space on the right will become the active basket and this can be reconfigured by right-clicking the name and selecting Properties. We can give the basket a title by right-clicking anywhere within the space and selecting the Text option. This will add a space where you can type in some text and format it using the options at the top of the window.

We can also add lists (see **LXF120** for more on this), group them together to make your basket management easier and add additional text elements too. Having added a few objects to the window, it's a good idea to practice moving them around, because what gets moved depends on where you grab an object. For instance, if you've grouped items together, you can remove an object from the group by hovering next to its name (the blue highlight will only cover that item) and dragging it away. This is also how you reorder items in a stack.

› Using the small minus icon, you can collapse a list down to just its title with a single click.

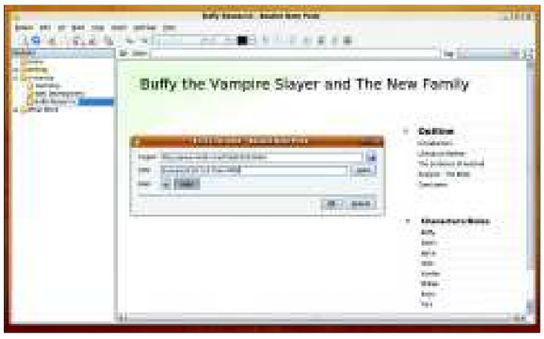

› It's *Basket*'s ability to work with web-based content and other apps that makes it so useful as a research tool.

Shift the mouse slightly to the left and the highlight should cover the entire group, which can then be dragged elsewhere.

The next thing we want to add is an image, which is done by dragging and dropping an image file from any folder into the active basket. When doing this, you'll be asked if you want to copy or move the picture. Move will delete the image from its original location, so you'll probably want to select Copy. Once in the basket, an image can be moved around just like anything else. We've created a title for the image and then grouped the two elements together, which can be done by either selecting both objects (click the first then Shift-click the second) and then clicking the paper clip icon in the toolbar, or by dragging one element so it attaches to the other (you'll see a thick black line appear beneath the original object). The second method means it's possible to, for instance, collapse the image beneath its title using the small minus icon. In this way it's possible to build up sophisticated stacks of content and then reduce them down to a title with a click.

Being visual

If you want to edit the image, you'll need to double-click on it to open it in a suitable editor. When you're done, your alterations are automatically saved to the *Baskets* version, which saves you importing the image twice.

We can also create associated links with our image (or a piece of text) by first selecting it – hover over it until the border appears and click it – then going to Insert > Link and providing the appropriate information. The link will become an appendage of the original text or image, meaning it will collapse using the minus icon as before. In our research basket, we've made a section called web links to house our collection of links associated with the project.

However, in-depth research relies on gathering information from a variety of sources, so look to the walkthrough for more on how *Baskets* integrates with other Linux apps and *Firefox*.

your research

Step by step: Gathering research

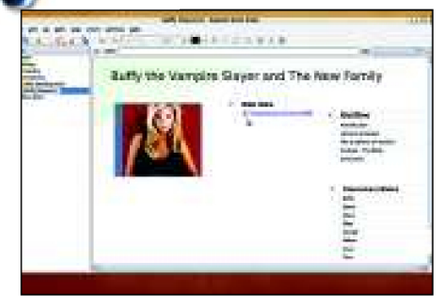

1 Drag and drop a link

To archive useful links you've found on the internet, navigate to the location you want to save, grab the favicon (the small icon next to the address bar) and drop it into your research basket. The link will then be formatted with the name of the page and can be moved around just like any other element within *Basket Note Pads*.

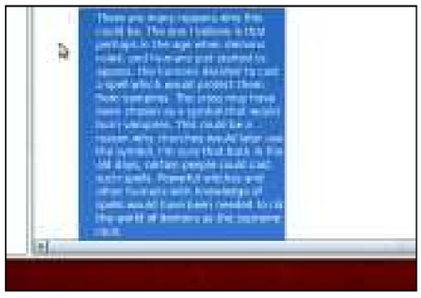

2 Take text from the web

If you find a particularly useful piece of text on the web, highlight and drag it from your web browser (or an application) into your basket, where it will appear as a new element. Note that this kind of text doesn't retain any information about its origin, so it might be best to also add a link to the original source. This text can be edited in the usual way.

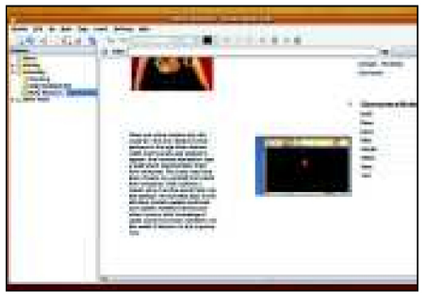

3 Take a screenshot

The problem with links is that they might expire or the text may change. A way around this is to take a screenshot of the site. Right-click a space in the basket and select Grab Screen Zone. *Basket* is minimised and the cursor turns into a crosshair. Draw around the page you want to save and it will be inserted into the basket.

4 Export a HTML page

HTML is a great way of backing up or sharing baskets with a friend or colleague who doesn't use *Basket*. Just navigate to File > Export > HTML Page and define a location for the file. This creates a single HTML file and a folder of the assets associated with it. Links in the basket will remain active, making this a good, if simple, web editing app.

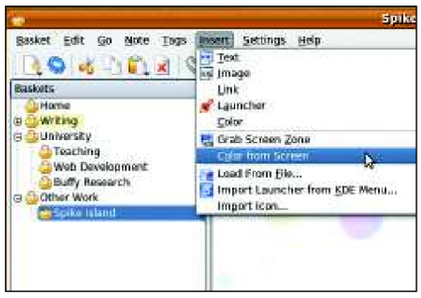

5 Build a colour palette

If you use this software to manage a design project instead, it's possible to build up a colour palette by sampling areas of the screen. So, you could open up a logo on your desktop in *Gimp*, select Insert > Colour From Screen and use the pointer to pick out a colour. The result is a swatch along with a hex value for the colour.

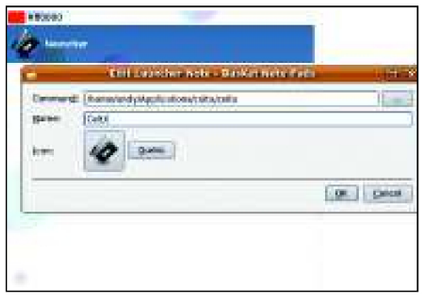

6 Create an app launcher

Finally, you may want to have applications associated with a piece of research – for instance, *CeltX* could be useful if you're writing a screenplay. To do this, right-click and select Launch. Either select your app from the list or browse to it using the ellipsis (...) icon. Now click OK to create a button that will launch your application. LXF

Next month Share more than just music and configure an effective firewall.

Gimp: Create a

Michael J Hammel continues his journey on the fickle wings of inspiration this month. He invites you to fire up Gimp and join him on his travels.

Our expert

Michael J Hammel is a contributor to the *Gimp* project and the author of three books on the subject, including his latest, *The Artist's Guide to Gimp Effects*.

After nearly 30 years of developing and using software, I sometimes find it hard to find motivation for my next project. The same holds true for my artwork, so when I come across something that reinvigorates my enthusiasm I tend to hold on to it like a CEO on to bailout funding. Enthusiasm is currency – and it must be spent quickly or the value declines.

So I must ask your forgiveness, dear reader, as I indulge my enthusiasm with another expert-level project based around inspiration. As with last month, this tutorial is more about trying to inspire budding *Gimp* users to dig around and combine ideas than revealing how to perform a specific process, although I'll cover the process as well.

Revitalise your burnt out creativity

This project, which I call Fire Goddess, is based on two online tutorials that I've merged together. The first of my inspirations, Mystic (**http://psdtuts.com/designing-tutorials/the-making-of-mystic**), produces in *Photoshop* what I can only describe as a Mayan god. This project is completely reproducible under *Gimp*, but I was unable to find suitable ornamental pieces to use here.

My second inspiration comes from a *Photoshop* tutorial that shows you how to create an image of a woman on fire **http://psd.tutsplus.com/tutorials/tutorials-effects/how-to-create-a-flaming-photo-manipulation**. This project is also reproducible in *Gimp*, but the lack of a warpable Transform tool makes the process more difficult – though not impossible – to complete.

The concepts behind the two tutorials are complementary and encourage experimentation. If you want to follow this amalgam project exactly, it requires a number of stock images that are either available for free or a low cost from stock image repositories (see *Take Stock, above*). It doesn't require the *Gimp Paint Shop* package I've described over the past few months, but having it will offer you more options.

I've created the project at 2350x2033, but if you need to to reduce memory and CPU requirements, scale the stock images down by 50% before starting.

Take stock

- **Our basis**
  Woman with hair blowing upward: **bigstockphoto.com/photo/view/1921777**
- **Ornate frames**
  Golden picture frame: **istockphoto.com/stock-photo-1710947-golden-picture-frame.php**
  Vintage gold frame: **istockphoto.com/stock-photo-4930364-vintage-gold-frame.php**
- **Flames**
  We used the fire pictures at:
  sxc.hu/browse.phtml?f=view&id=1093986
  sxc.hu/browse.phtml?f=view&id=1099709
  sxc.hu/browse.phtml?f=view&id=1104673
  sxc.hu/browse.phtml?f=view&id=1143736

This flaming hot beauty is the result of all your hard work.

» **Last month** We destroyed a city with Gimp and a graphics tablet.

fire goddess

Step by step: Set the world alight!

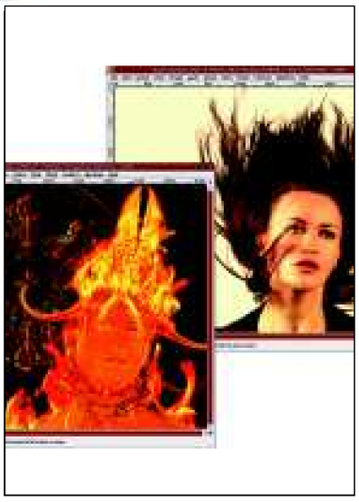

1 A simple process

Let's start by looking at the original image and the final image. As you can see, plenty has changed, but there are two notable alterations: firstly, I've removed the hair of the girl and, secondly, the image is flame coloured.

The reason our goddess is hairless is that I'd originally intended this project to closely follow the tutorial for creating a flaming girl. With that in mind, I chose a base image with hair that should be easy to convert to flames using *Gimp*. However, because that process used a feature of *Photoshop* that isn't easily duplicated in *Gimp* (specifically, the warpable Transform tool), I changed this project, making the hair in the source image irrelevant.

Throughout the project you'll also notice that flames are not just components of the image, but also used for colouring. The easy way of using a flame to colour a layer is to desaturate the layer to be coloured, add a flame layer above it and set the flame's Layer Mode appropriately. However, if you place anything below that layer, the black region around the flame may block it from being visible. To avoid this, we'll use layer masks on flame layers or simply copy and paste flames (without the black backgrounds) from their source image.

The final process for creating the Fire Goddess can be reduced to these basic steps:

1. Isolate the girl from the background.
2. Detect the edges of her head and create multiple corresponding layers to get a perfect selection.
3. Blend in ornamental elements.
4. Do some fine blending to remove excess details.
5. Colour the image with flames.
6. Create the flaming hat.
7. Add the background emblem.

Note that even if you follow this tutorial exactly, you won't produce the same image as me. The steps I'll present allow for too many variations, especially when creating the flaming hat. Don't be discouraged if your first attempt is less than ideal, though. I went through seven different designs before settling on this process and the first four or five of those are – to be kind to myself – not something I'd ever show in public.

2 Isolate the girl

The first step in this project is to isolate the girl from her background. We'll do this primarily because the background in the source image for the girl isn't required and its colour may cause problems when mixed with other layers. Removing it will also enable me to add background elements behind the girl's head.

To isolate the girl, first Desaturate (Colours > Desaturate) the source image with the Luminosity setting selected. The reason for using this setting is that the source image has a yellowish tint when converted from its original colour space (provided in the image when downloaded from the stock image website) to the colour-managed display in *Gimp*. I say colour-managed, because I configured the monitor profile for my Acer X203W monitors in *Gimp*'s Preferences. Thankfully, the yellow tint converts to a bright white when the Luminosity option is selected in the Desaturate window.

Now I can use Select By Colour from the *Gimp* toolbox and click on the white background. With the Tool Options Threshold set to 50, this will select all of the background, a little of the hair and some of the face and neck. To avoid selecting any skin, switch to Quick Mask mode and paint out the white areas on the girl with a large black brush. Now switch off Quick Mask mode and Invert the selection. This selection is then copied and pasted to a new layer, which I name Girl. I then fit the new layer to the canvas using Layer > Layer To Image Size so that any modifications I make won't be restricted by the size of the layer.

3 Cutting her hair

In the original layer (which is titled Background by default), I use Ctrl+A to select the entire layer and then drag the black foreground colour into the canvas to fill that layer with black.

Now it's time to get rid of the girl's hair. Just like most operations that involve components of an image, I don't really want to remove these from the layer itself, but rather leave them in place and mask them with a layer mask. The reason for leaving these elements in place is so that I can decide later to reincorporate the component (in this case hair) into the project. This kind of flexibility is a real assets of editing images digitally, so I encourage you to make use of it.

In this case, I add a white layer mask to the image. Starting with a large brush and switching to smaller brushes as I proceed, I paint with black on top of the girl's hair in the layer mask until most of the hair is removed. I leave a small amount of hair because I will be using edge-detect filters later and I know the lines in the hair will provide another visual feature using those filters. Afterwards, I switch off the Girl layer's visibility so it won't interfere with the next steps. »

4 Edge detection

The next step in this project is to create detailed outlines of the girl using two edge-detection filters: Edge (Filters > Edge Detect > Edge) and Neon (found by going to Filters > Edge Detect > Neon). These produce similar results that I'll combine later with Layer Modes.

To begin with, I duplicate the Girl layer and name the duplicate layer Edge Detect. After that, I then open up the Edge Detect filter and select the Prewitt Compass algorithm. I choose an algorithm based on the preview and here I selected the one that showed the most detail. Apply this filter to the layer and then set the Layer Mode to Screen.

The Girl layer is duplicated again, this time naming the new layer Neon Detect. This layer is moved to the top of the layer stack and the colours are inverted in this (Colours > Invert) before applying the Neon edge detection filter. The Layer Mode is set to Screen and, finally, this layer is duplicated once, resulting in three layers dedicated to edge detection stacked above the Girl layer.

Zoom in now and you may find some strands of hair that weren't masked in the last step. These can be left, or you can remove them with the Eraser. If you do this, the Eraser will probably need to be applied to the same area on each of the three edge detection layers. These would be too difficult to see in print, so I'm not showing them here. In my project, however, I erased the extra strands of hair – I'm obnoxiously detail-oriented.

5 Blend in the ornamental elements

Now it's time to get our goddess looking ornamental. The best place to find suitable designs to use for this are picture frames and building architecture. This project has stuck with picture frames only because I couldn't find any really good photos of Victorian architecture.

There are three ornamental pieces I'm adding to this design: one on the forehead, one at the junction of head and neck, and one to form a collar. The first two come from a rectangular frame while the latter comes from an oval frame.

I start with a rectangular frame and Desaturate it. Because the frame is fully enclosed in the solid-coloured background and also encloses a solid-coloured area, selecting the frame is easiest using the Fuzzy Select tool set with a Threshold of 50. Clicking once outside the frame and again inside it creates a selection of the background. Grow this by one pixel (something I nearly always do to avoid leaving light-coloured pixels along the edge of the selected object) and then Invert the selection. Copy and Paste the frame into the working canvas as a new layer at the top of the layer stack. Name this layer Ornamental:Forehead and ensure the Layer Mode is set to Normal.

Once that's done, Scale, Rotate and position the Forehead layer over the girl's face. Use a layer mask to remove the left and right sides and bottom edge of the frame. Duplicate the layer and set the duplicated layer's Mode to Dodge. Dodging the copy over a desaturated Normal Mode original enhances contrast (which makes it look shinier).

6 Continue decorating the goddess

The ornament around the neck is created in exactly the same way, albeit with a new copy of the frame. A new copy is required because the angle of rotation, size and position will be different to the Forehead version, not to mention being masked differently. In this layer, the bottom edge of the frame is kept while the rest is masked out.

Similarly, the collar is decorated with the oval frame. This time the frame is flipped vertically before rotation, sizing and positioning. Finally, a mask is added to keep just the lower section of the frame.

Note that the frame images were significantly smaller than the main image window and the girl. This required increasing the size of the frames. Scaling an image up is problematic because the scaled image quickly becomes pixellated. However, in this case the pixellation is hardly noticeable because of the weakly identifiable shapes (the ornate structure has no discernible pattern). Using blending layer modes also hides some of the pixellation and our colouring technique hides it even more. Large prints, however, would likely show the pixellation. Therefore, if you intend to make any prints of this image, they should to be no larger than the size of a magazine page. (Pretty convenient, eh?)

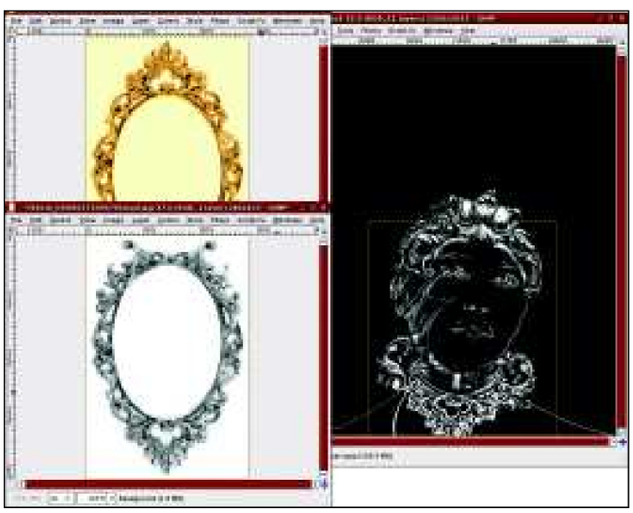

7 Blending to remove excess details

The ornaments didn't cover as much of the girl's hair as I'd originally thought it would, so to get rid of these unsightly blemishes on our goddess, it's time to turn our attention to the Eraser and Smudge tools.

Smudging in the edge detect layers is possible (white and black smudged together produce shades of gray), but I'd rather leave some of the line details in the face. To accomplish that, I'll use the Eraser with a Grunge brush and set the Brush Dynamics to increase Opacity, Hardness and Size when pressure is applied using my Wacom pad. It's worth mentioning here that throughout this project I haven't needed to break out the pad, but this is one place where it comes in handy. If you don't have a graphics tablet to hand, you can just use a smaller brush and vary the Opacity manually.

Now I apply the Eraser to each edge-detect layer in turn, starting with the upper layers. I do only a small amount on each layer before moving to the next, then rotate back to the top and repeat. My goal is to try and keep as much of the detail as possible while losing the obvious hair lines.

8 Smudging out excess hair

Removing the hair also means Girl layer comes back into view. Now switch off the visibility of the three edge detection-related layers and turn on the visibility of the Girl layer. Enable the Layer Lock for the Girl layer and then use the Smudge tool with any appropriate brush to blend the hair out of this layer. Try to keep the lighting smooth across the cheek by clicking in lighter areas first (the cheek) and dragging into darker areas (the hair). Smudge the neck and clothing as well.

With the Layer Lock enabled, you don't have to avoid smudging outside the girl's outline. While this is something I desired for this project, you could express your creativity by disabling the Lock and smudging all over the place. Just be sure to colour the smudged areas with some fire later. When you're done, make the edge detection layers visible again.

9 Add flame colouring

Colouring is the next step. The girl is coloured first, simply by using Colourise with Hue, Saturation and Lightness set to 30, 90 and -30, respectively. Next, a roaring fire image (flame and black background) is copied in as a new layer to the top of the layer stack. It's scaled so that the flame fits completely over the girl, then the layer is cropped to fit the image (Layer > Layer to Image Size). Set the Layer Mode to Hard Light.

When Layer To Image Size is used, the layer may not span the width of the image and you'll see a white band. If this happens, click the white band with the Fuzzy Select tool, Grow by it one pixel and fill the selection with black.

At this point, the flames obscure the girl's face and you have two options: leave them or use the Smudge tool on the flame layer to let the face and ornaments show through. I chose the latter option.

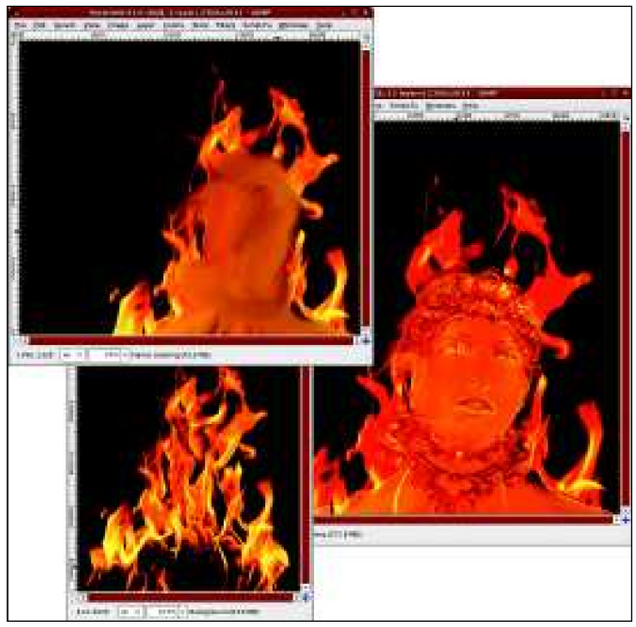

10 Create the flaming hat

Finally, we're at the really fun and creative part. The process for creating the fire hat is remarkably simple, but it will more than likely take you a couple of attempts to get it just right. Start with a flame image like the one shown below – a flame, not a roaring fire. Copy just the flame (without the black background) into the project's window. Name the new layer Flame Hat. Scale this layer to fit from the girl's forehead and stretch it to the top of the image.

Duplicate this layer, flip it horizontally and position it close to the original flame layer. Use a mask to merge the flames together if necessary. Then merge the two layers by selecting the top layer of the two and selecting Layer > Merge Down. Move the merged layer to the middle of the image window and size the layer to the image (again using Layer > Layer to Image Size). Now you have the basic flame hat frame in place with some space around it to create the flame swirls.

Use the iWarp filter to create the flame swirls. Now, here's the tricky part. The left side of the hat will be swirled using the Swirl CW mode and the right side of the hat will be swirled using the Swirl CCW mode. For this particular flame, I found two spots on each side which would make good swirls. Clicking (never dragging) at various points in the preview, I was able to create two swirls with a nearly uniform appearance. Clicking and holding also works, but you'll need to hold for a very short amount of time or the swirl gets too big.

You should experiment with the Deform Radius and Deform Amount while you do this, as well as where to click in the preview. If the preview gets messed up, just hit Reset and try again. It may take patience, but if you experiment enough, you'll get it in the end.

After applying the iWarp settings to the Flame Hat layer, position and Rotate the hat in place over the girl's head. Set the Layer Mode to Screen, then add another copy of the roaring fire above this layer, scaled so the whole flame fits inside the hat. Set the Layer Mode for the roaring fire layer to screen and use a layer mask to blend it into the hat.

11 Add Background Emblem

The last step is extremely simple. To add a sort of Mayan feel to this project, I selected the Mandala brush from the *Gimp Paint Shop* collection. With the Paintbrush tool selected, I increased the scale of the brush until a single brush click was roughly three quarters of the size of the image window. Then I added a transparent layer above the background, made the foreground colour white and clicked once in this new layer.

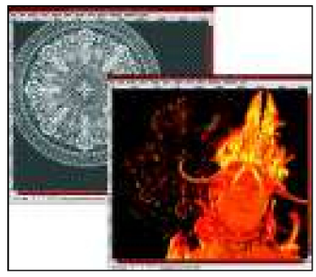

Above the Mandala imprint, I once again added the roaring fire image in a new layer. Scale the fire to cover most of the imprint and then set the Layer Mode to Multiply. And that's it.

Summary

Earlier I mentioned that I toyed with the idea of removing the girl from the image. This didn't look right until I added the Mandala imprint layer and coloured it with flame. Now when I remove the girl layer, the imprint shows through, which looks pretty good. Why not take this basis and experiment some more? Soon you'll have a great variation on this theme. LXF

» Next month Fire up the dilithium crystals and add a touch of warp speed!

OOo: Automate

Marco Fioretti reveals how to create a slew of OpenDocument invoices and multiple choice tests automatically to save you from monotonous chores.

Our expert

Marco Fioretti Is the author of *The Family Guide to Digital Freedom*. He's also a free software activist and programmer.

We often use computers to produce many different versions of a single document, but updating our files manually only makes sense if our changes are infrequent. However, there are many ways to avoid wasting valuable seconds – as we demonstrated in **LXF119** when we looked at how to automatically update spreadsheets. This month, we'll look at the big picture: how to go from raw data to any kind of productivity document with minimal manual work. We'll also manage to avoid opening *OpenOffice.org* (well, almost), thanks to the characteristics of its default file format: OpenDocument or ODF.

An ODF file is just a zip archive, with pictures and macros in their own folders and the actual text – written in XML format – inside a file called **content.xml**. To create a new ODF file, you must open the archive, change the text inside **content.xml**, put any new pictures in their folder and zip everything up again. Here, we'll show you how to do this with ODF text documents and really simple shell scripts.

An ODF multiple choice test system

Let's begin with a practical problem: generating a formatted multiple choice test with randomly generated questions. Let's assume for simplicity's sake that the test only has one question, and it's in a document that's laid out with a test name, test description, question and three choices, each with its own picture and caption. This file, which we'll call **sample_multiple_choice.odt**, is our starting template. In order to use that template to generate our test in OpenDocument format, first we have to write down all the strings we want in the test in an ASCII file (we've called this file **my_test_data_1**) as follows:

```
marco => cat my_test_data_1
TEST_NAME='Gnu/Linux Compatibility test'
DESCRIPTION='This is an highly scientific test aimed to discover your most hidden feelings about Free Software'
QUESTION_NAME='What is your preferred Gnu/Linux distribution?'
FIRST_CAPTION='Is it Ubuntu?'
SECOND_CAPTION='...or Mandriva?'
THIRD_CAPTION="Or maybe you're a Fedora fan?"
```

While the format of this file is simple, there are a couple of things that you shouldn't overlook: to begin with, this is perfectly valid shell syntax. Every line of the file assigns a value to some shell variable that will be used in the main script. The second is that although there's no general rule, you must be careful with quotation marks here. For example, the last line has to use double quotes, because it already contains a single quote character (in **you're**).

Picture preparation

Besides the strings, you also have to prepare the three pictures in the same order they will appear in the document. In our example, they're found in the **Linux_Test_Pictures** folder and we've called them **01_ubuntu_logo.png**, **02_**

Using ODF means you can make documents without using *OpenOffice.org*, but you'll still need *OOo* to create the initial template.

» **Last month** We discovered how to generate automatic slideshows.

your work

OpenDocument files have no secrets: once you unzip them, the text, metadata and pictures are all in plain sight.

mandriva_logo.png and **03_fedora_logo.png**, respectively. We found that numbering them was necessary to guarantee they're used in the right order. Once everything is in place, run the following command

```
marco => test_generator.sh sample_multiple_choice.odt my_test_data_1  Linux_Test_Pictures/*
```

to finish your first test.

Reuse the script

The best part of this process is that the method remains generic as long as the layout is constant. As proof of this, place three appropriate pictures (the first of mountains, the second of balmy beaches and the last of a historic city) in the **Holiday_Pictures** folder and the add following text in an ASCII file called **my_holiday_data_1**:

```
marco => cat my_holyday_data_1
TEST_NAME='Holiday orientation test'
DESCRIPTION='How to discover which kind of vacation is better for you'
QUESTION_NAME='Which of these places would you like to visit first?'
FIRST_CAPTION='Wild mountains?'
SECOND_CAPTION='Sunny tropical beaches?'
THIRD_CAPTION="Or historic cities?"
```

Then finish the process by running:

```
marco => test_generator.sh sample_multiple_choice.odt my_holyday_data_1  Holiday_Pictures/*
```

Isn't it great? Now all you'd have to do to generate 1,000 similar files would be to run the script below in a loop.

Here's the result of running your simple OpenDocument processing script: your initial template is filled with dynamic data, all without needing *OpenOffice.org*.

Listing 1: test_generator.sh

```
 1  WORK_DIR=odt_test_generator_temp_dir
 2
 3  rm -rf $WORK_DIR
 4  mkdir  $WORK_DIR
 5  FILENAME=`basename $1 .odt`
 6
 7  cp     $1 $WORK_DIR/my_template.odt
 8  cp     $2 $WORK_DIR/my_data.sh
 9  shift  # remove $1 from argument list
10  shift  # remove $2 from argument list
11
12  ## copy all pictures into work directory
13  touch $WORK_DIR/new_pictures_list
14  for VAR in "$@"
15  do
16  CURRENT_FIG=`basename $VAR`
17  cp $VAR $WORK_DIR/
18  echo "cp ../$CURRENT_FIG  " >> $WORK_DIR/new_pictures_list
19  done
20
21  ## preparation
22  cd     $WORK_DIR
23  mkdir  work
24  mv     my_template.odt work
25  cd     work
26  source ../my_data.sh
27  unzip  my_template.odt > /dev/null
28  rm     my_template.odt
29
30  ## replace text strings
31  sed "s|Test Name|$TEST_NAME|"  content.xml \
32  | sed "s|Question|$QUESTION_NAME|"        \
33  | sed "s|Test description and general conditions|$DESCRIPTION|"        \
34  | sed "s|first caption|$FIRST_CAPTION|"        \
35  | sed "s|second caption|$SECOND_CAPTION|"        \
36  | sed "s|third caption|$THIRD_CAPTION|"        \
37  > custom_content.xml
38  mv custom_content.xml content.xml
39
40  ## get names of embedded pictures and replace them
```

```
41  tr ">" "\012" < content.xml | grep 'draw:image
xlink:href' | cut '-d"' -f2 > ../pictures_list
42  paste ../new_pictures_list ../pictures_list > ../copy_
pictures
43  source ../copy_pictures
44
45  ## zip everything, rename it as .odt file and clean up
46  find . -type f -print0 | xargs -0 zip ../$FILENAME > /dev/
null
47  cd ..
48  mv $FILENAME.zip ../new_$FILENAME.odt
49  cd ..
50  rm -rf $WORK_DIR
```

The **test_generator.sh** script takes, as arguments, the ODF text template, the ASCII file containing text strings and all the pictures that should go in the new document. The first 10 lines define a temporary directory and copy all the necessary files into it. Meanwhile, the **shift** commands in lines 9 and 10 remove the ODF template and the data from the argument list variable **$@**. This is necessary to make the loop in lines 14 to 19 work only on the graphics files. It makes a local copy of each picture, but its most important part is line 18, which creates a **new_pictures_list** file that looks like:

```
marco => cat new_pictures_list
cp ../01_ubuntu_logo.png
cp ../02_mandriva_logo.png
cp ../03_fedora_logo.png
```

You'll discover why we need to generate a file like that in a second. For the moment, let's go back to the code. Lines 22 to 28 move us over to the **work** directory, load all the variables set in the data file (line 26), unzip the template and then, finally, we can begin generating the new ODF file.

Pay attention to the text string replacement section (lines 30 to 38), comprised of a series of **sed** commands connected in a pipe. Each invocation of **sed** replaces one string of placeholder text in the **content.xml** file with the content of one of the variables set in the **my_test_data_1** file. If you want to customise this script to create other documents later, it's this section of code and the data file you'll need to change.

After the new **content.xml** file has been created, we then replace the original pictures with the ones we want to appear in the template. This happens in two steps. Line 41 uses the **tr** and **grep** commands to extract and write the names of all pictures to the **pictures_list** file. These are written in xlink XML attributes inside **content.xml**, and the process should produce results similar to:

```
marco => cat pictures_list
Pictures/100000000000008F0000008A1DC84E9A.png
Pictures/1000000000000009100000008EDC1B9151.png
Pictures/1000000000000008100000008C06E4E423.png
```

After that, we just need to paste **new_pictures_list** and **picture_list** together line by line (lines 42 to 43), to produce the command file sourced in line 43:

```
marco => cat copy_pictures
cp ../01_ubuntu_logo.png Pictures/100000000000008F00
00008A1DC84E9A.png
cp ../02_mandriva_logo.png Pictures/10000000000000910
000008EDC1B9151.png
cp ../03_fedora_logo.png Pictures/1000000000000081000
0008C06E4E423.png
```

All the pieces are now ready. What happens in the rest of the script is simply zipping everything up and renaming the results with an ODT extension. As a second example, let's now look at how to generate invoices in ODF format. There are articles on how to do this already, but they all require you to start and use *OOo* manually. What we want is to have our computer do all the work for us. The procedure is the same as we just discussed, only simpler due to the absence of pictures. First, set up a template and then create this data file:

```
marco => cat my_invoice_data_file
INVOICE_DATE='2009/03/20'
VENDOR_CODE='007'
PO_NUMBER='LXF 10541'
TOTAL=100
ISSUE=150
DESCRIPTION='Nothing less than the best tutorial ever
published by Linux Format!'
```

to transform the template into an invoice.

Now run the script below, which is derived from the **test_generator.sh** shell script we looked at in the first example:

Listing 2:

```
1  WORK_DIR=odt_invoice_generator_temp_dir
2
3  rm -rf $WORK_DIR
4  mkdir  $WORK_DIR
5  FILENAME=`basename $1 .odt`
6
7  cp     $1 $WORK_DIR/my_template.odt
8  cp     $2 $WORK_DIR/my_data.sh
9
10 ## preparation
11 cd     $WORK_DIR
12 mkdir  work
13 mv     my_template.odt work
14 cd     work
15 source ../my_data.sh
16 unzip  my_template.odt > /dev/null
17 rm     my_template.odt
18
19 ## replace text strings
20 sed "s|__INVOICE_DATE|$INVOICE_DATE|"  content.
xml \
21 | sed "s|__VENDOR_CODE|$VENDOR_CODE|"      \
22 | sed "s|__PO_NUMBER|$PO_NUMBER|"          \
23 | sed "s|__TOTAL|$TOTAL|g"                 \
24 | sed "s|__ISSUE_NUMBER|$ISSUE|"          \
25 | sed "s|__DESCRIPTION|$DESCRIPTION|"      \
```

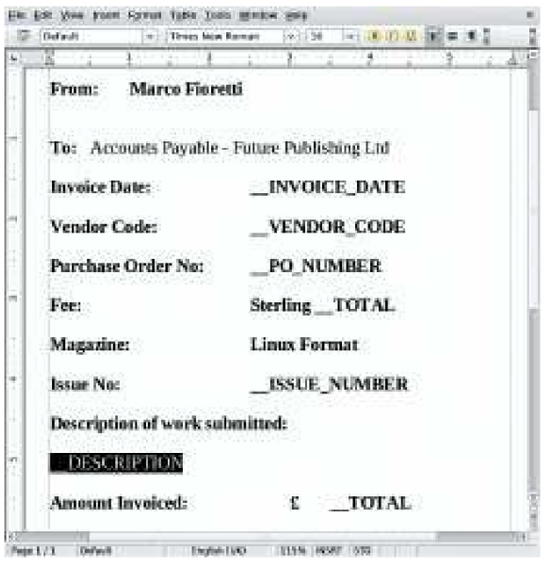

Different need, same solution: one simple template is all you need to get your computer generating invoices.

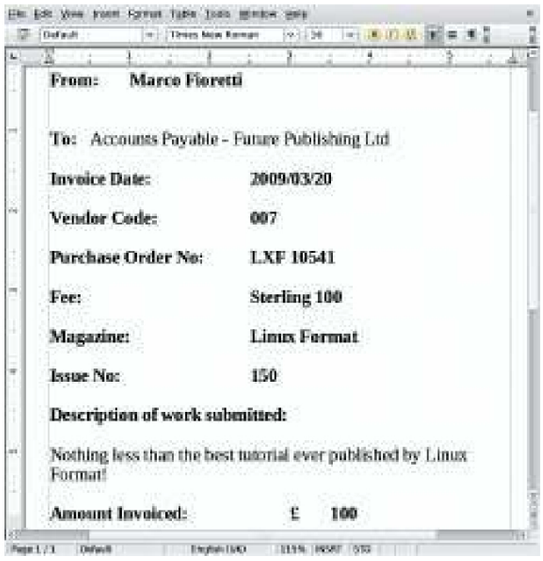

From: Marco Fioretti

To: Accounts Payable - Future Publishing Ltd

Invoice Date:	2009/03/20
Vendor Code:	007
Purchase Order No:	LXF 10541
Fee:	Sterling 100
Magazine:	Linux Format
Issue No:	150

Description of work submitted:

Nothing less than the best tutorial ever published by Linux Format!

Amount Invoiced: £ 100

› Here's the result of your labour: a final invoice, ready to be printed or sent via email.

```
26 > custom_content.xml
27 mv custom_content.xml content.xml
28
29 ## zip everything, rename it as .odt file and clean up
30 find . -type f -print0 | xargs -0 zip ../$FILENAME > /dev/
null
31 cd ..
32 mv $FILENAME.zip ../new_$FILENAME.odt
33 cd ..
34 rm -rf $WORK_DIR
```

Much of this should be familiar to you from what we covered in Listing 1, but take a moment to look carefully at lines 19 to 25. As we've already mentioned, if you change the template, you must add one **sed** command for each string substituted. If the same string occurs multiple times, as it does for total price here, remember to add the **g** (global) option to **sed**, otherwise it will only see the first occurrence of that string (see line 23).

What about MS Office versions?

In a perfect world, everybody would be already using OpenDocument in a paperless office and in such a paradise there would be no need to bother about the file formats other people are using, or the licence they run to open files. However, until that day arrives (if it ever does), we can't live without a way to print files or convert them for those poor souls still chained to applications that are only capable of understanding *Microsoft Office* formats.

Thankfully, both printing and conversion can be done automatically. Unlike the generation of the ODF files, however, this is a step that needs *OpenOffice.org*. You'll need to add one line to your scripts that launches OpenOffice without the GUI, then execute an *OOo* macro to convert the ODF file to PDF (for printing) or DOC format. There are several macros to use for this purpose – two of the best are *SaveAsPDF* and *SaveAsDoc* from **www.xml.com/pub/a/2006/01/11/from-microsoft-to-openoffice.html**. Alternatively, there's a PDF generation macro at **http://linux.derkeiler.com/Mailing-Lists/Fedora/2008-06/msg00561.html**.

To print the PDF file from your script you can just feed it to the **lpr** command. Using *OpenOffice.org* from the command line is explained at **http://tinyurl.com/rybr9d**, but this is the correct syntax:

```
soffice -invisible macro://path-to-macro($FILE)
```

The **-invisible** option is what makes *OpenOffice.org* start without the GUI. The file to process must then be passed as an argument (**$FILE**) to the macro.

What have we learned?

In this tutorial, we've explored a method to automate repetitive tasks with documents, which has several important advantages. To begin with, it works without having to run *OpenOffice.org* (with the exception of printing), meaning it could run on a server. Also, this method doesn't rely on any sort of relational database and, above all, it's simple! There are many cases when XML-based tools (such as those mentioned in *Resources, below*) are too much trouble to study and deploy, but using the tricks here, you can at least avoid wasting time on repetitive edits. In fact, anybody with just a basic grasp of shell scripting can generate, modify or analyse hundreds of ODF text documents with just a few minutes of basic coding.

But it's more than just ease of use that makes our process attractive. For starters, the fundamentals of the process may well be more familiar ground for you. We reckon most Linux users are more comfortable with *Bash* scripting than with StarBasic – the language *OOo* macros are written in. This links in to our final reason for bypassing *OOo* macros: implementing all the decision logic outside *OpenOffice.org* gives you much more power, flexibility and potential for integration with other tools, from *Cron* jobs to mass mailing and graphics processing. It would be simple, for example, to add the *ImageMagick* tools for scaling and framing pictures to our first script before adding each one to our ODF document!

The two scripts presented here aren't perfect, of course. They aren't particularly flexible or scalable, nor are they all that robust in their current form. For one thing, there's no error management: if one of the files is missing, or if you feed more pictures to the test generator than there are in the original template, things will go wrong without leaving any useful indication of why.

However, although you can refine them, the purpose of this series isn't to develop industrial-strength solutions. First of all, this tutorial proves that OpenDocument has no secret depths: an ODF file is a normal zip archive containing plain text and pictures that are compatible with all kinds of software. Store your data in ODF format and it's easy to recover all that information, even if *OpenOffice.org* were to disappear tomorrow. The second, and most important, thing we want to show you is that ODF is quick and easy to hack to suit your needs, and ripe for experimentation. LXF

Resources

If you want to get really serious about ODF processing, the tricks in this tutorial won't be enough to do all you need. If that's the case, you'll have to either use XML-based tools such as *Odfpy* (**http://odfpy.forge.osor.eu**) or check out the book *OpenDocument Essentials* by J David Eisenberg, which you can buy at **www.lulu.com/content/207835** or read online at **http://develop.opendocumentfellowship.com/book**.

» Next month We'll look at processing the data we've coralled into ODF files.

Mail server: Get

Part 8: Complete your mail server setup with Postfix. **Neil Bothwick** shows you how, and teaches you how to filter out spam and viruses in the process.

› ***Postfix*'s main configuration file is long, but heavily commented and filled with sensible defaults.**

Our expert

Neil Bothwick Neil has a computer in every room, but for security's sake won't disclose the location of his central server.

Back in **LXF115**, we set up an IMAP mail server to provide access to your mail from anywhere. Last month, we added webmail capabilities to that server, but we still currently rely on having mail delivered to your ISP or another mail service first and then pulling it from there with *Fetchmail*. Now it's time to complete the circle by adding a Simple Mail Transport Protocol (SMTP) server, which will enable you to have mail delivered directly to your IMAP server. It also means you can set up some useful virus and spam filtering before the mail is delivered. In essence, SMTP sets the rules that govern how emails should be transferred between computers.

The main component in this process is the Mail Transport Agent (MTA), which is responsible for receiving mails from one computer and sending them to another. You may also see the abbreviations MUA and MDA. An MUA (Mail User Agent) is the program you use to manage your mails, while an MDA (Mail Delivery Agent) is responsible for taking the mail from a server and delivering it to a user's mailboxes. Most MTAs can act as a basic MDA, but it's common to use a separate, more capable program. We used *Procmail* as our MDA when collecting mail with *Fetchmail*, so we'll continue using it here.

Choosing to use a separate MDA has another advantage: you can daisy-chain programs. Instead of the MTA passing the mail directly to the MDA, it can pass it to a spam filter, then a virus checker and finally over to *Procmail* for delivery. This means that before any mail arrives it has been thoroughly checked for unwelcome content and either flagged as undesirable or filtered somewhere else.

There are several popular mail servers you can use with Linux, including *Postfix*, *Sendmail* and *Exim*. *Sendmail* and *Exim* are both capable programs – even if *Sendmail*'s configuration file syntax is so complex it could baffle the SETI project – but ultimately we've plumped for version 2.5.6 of *Postfix*, because it was released at the beginning of this year and should be readily available. If you don't have it installed already, you can add it from your package manager in the usual way.

Setting up Postfix

Postfix has several large, well-commented configuration files found in **/etc/postfix**, but the default setup needs little alteration. The main config file is **/etc/postfix/main.cf** and you should read the comments in here before changing anything. Now, assuming your domain is **mydomain.co.uk** and you also want to handle mail for **mydomain.com** and **example.com**, the changes you should make begin with:

```
proxy_interfaces = 192.168.1.1 # optional
```

Changing **proxy_interfaces** may not be necessary, but it can help to guard against mail loops. Omit it if you have a direct connection without a router or NAT, or if you do include the change it should be set to the IP address of your internet gateway. Now set:

```
myhostname = mydomain.co.uk
```

This is your hostname, and the hostname that your server will use when connecting to others. It must resolve to your

» **Last month** We added some web applications to our Apache server.

rid of spam

external IP address, otherwise some servers may refuse the connection as potential spam. Next, you'll need to alter **mydestination** as follows:

```
mydestination = mydomain.co.uk,mydomain.com,example.com
```

By doing this, you've created a comma-separated list of the domains that this server will handle email for – when it receives mail for a domain not in this list, it will try to send that mail on to a server that can handle it. Now change **myorigin** to point towards your domain address:

```
myorigin = mydomain.co.uk
```

Note that all mail from this server will appear to come from this domain now. It will also be appended to any addresses with no domain part (local mail).

If mail arrives for an invalid user, it would normally be rejected. The **luser_relay** parameter specifies a default recipient for these mails. If you want to be able to receive mails for any username, valid or otherwise, set this to:

```
luser_relay = username # optional
```

This enables you to use different usernames for different purposes, but can increase the amount of spam you'll receive, as a lot of junk mail is sent to random usernames.

When you send email, *Postfix* normally tries to connect directly to the recipient's mail server. If your ISP doesn't allow this, or your ISP's range of user IP addresses is on a spam blacklist, you can use the following to have all mail sent through your ISP's mail server:

```
relayhost = smtp.myisp.co.uk # optional
```

You probably have this set up as the SMTP server in your mail client at the moment. Finally, you should alter:

```
mailbox_command = /usr/bin/procmail
```

Postfix is able to deliver mail to the users' mailboxes directly using either the **home_mailbox** or **mail_spool_directory** setting, but it can also call an external delivery agent, which provides a lot more flexibility. Since we've set up *Procmail* to work with *Fetchmail* in **LXF115**, we will stick with that (for now at least). If you don't have **LXF115** to hand, put the following into **/etc/procmailrc** as a starting point for *Procmail:*

```
MAILDIR=/var/spool/mail
DEFAULT=$MAILDIR/$LOGNAME/
LOGFILE=/var/log/procmail
```

Setting up your router

In order to receive mail from outside, you need a static IP address, which many ISPs will provide (although some charge extra for one). You also need a DNS MX record for your domain that tells other mail servers where to deliver mail for you. Wherever you registered your domain name should be able to set this up for you. SMTP works through port 25, so you will need to open this on your router and forward it to your server.

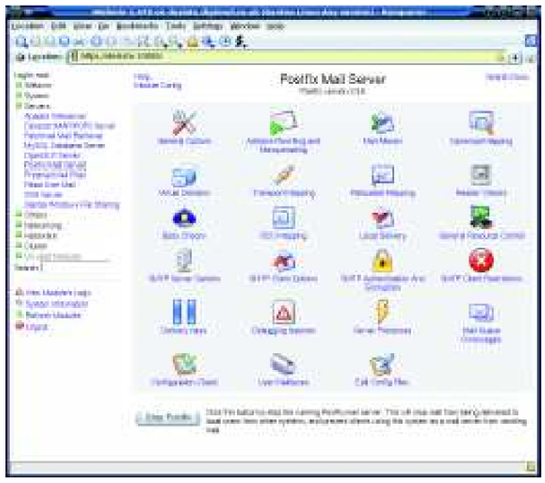

These aren't *Postfix* config options, instead they're the categories for those options: there's a lot of choice if you need it.

With that done, start (or restart) the *Postfix* service, set your mail client to use localhost (or the hostname of the computer running *Postfix* if it's on a different computer on your network) on port 25 and send yourself a test mail. Before you hit send, run

```
sudo tail -f /var/log/messages | grep postfix
```

in a terminal to output all new data written to the system log file and filter anything relating to *Postfix*. When you send a mail, you should see a few lines of output showing that a connection was established from your mail client, that a mail was received and that it was delivered to *Procmail*. If you see an error here, it should refer to a configuration setting that you can correct. If there is no error and the mail is not delivered, it was probably successfully passed to *Procmail*, but the delivery failed. Check **/var/log/procmail** for details.

Authorised access

In the old days, when the only inhabitants of cyberspace were academics and geeks, mail servers worked in a friendly, co-operative manner. Any server would forward any mail it received to the correct recipient, or a server closer to them. All that stopped when Sanford 'Spamford' Wallace and his colleagues realised that this was an easy way of sending out huge amounts of unsolicited mail. Now mail servers are set up with relaying disabled. This means that a server should only accept and deliver a mail if it meets one of three criteria:

- The email is intended for one of the recipients on the server's domain list.
- The mail is sent from a machine on the server's domain list.
- The user sending the mail authenticated themselves before the email was sent.

The first two are the default settings for any sensible MTA, meaning that mail is only accepted if the destination or source are on the server's network. This is why you may find yourself having to change the SMTP address in your mail »

Quick tip

These programs are daemons and don't normally provide feedback to the terminal that ran them. All useful output goes to a log file, either their own or via *syslog*. Use **tail -f** to watch the log file in real time while testing.

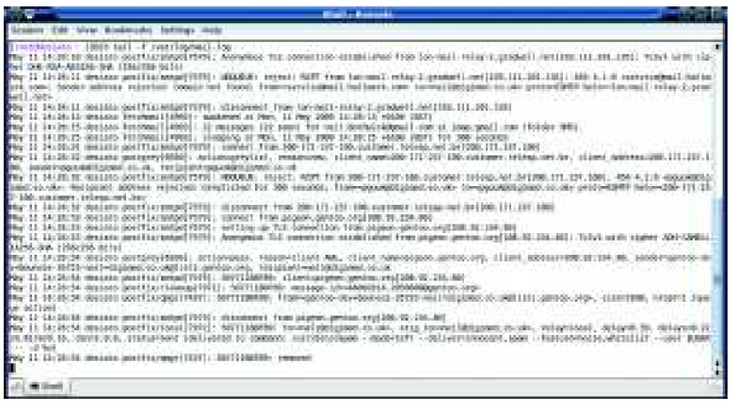

Keep an eye on the log files when testing new options – you'll find errors are usually obvious.

client when you switch from one network to another, because you're forced to use the server for that network. However, authentication avoids this problem, since authenticating means the user is explicitly trusted and the server will forward mail for them. You'll need this on your own server if you want to be able to use it to send mail, whether you're on the local network, using a wireless connection in a coffee shop or connected via a 3G mobile dongle. *Postfix* uses *Cyrus SASL* for authentication, so make sure this is installed and that the *saslauthd* service is started when you boot. Now add these lines to the end of the **/etc/postfix/main.cf**:

```
smtpd_sasl_auth_enable = yes
smtpd_sasl_security_options = noanonymous
smtpd_sasl_local_domain = $myhostname
broken_sasl_auth_clients = yes # optional
smtpd_recipient_restrictions = permit_sasl_
authenticated,permit_mynetworks,check_relay_domains
```

The line marked as optional is required only with some versions of *Outlook Express* and *Microsoft Exchange*. After reloading the config, you should still be able to send mail from your mail client without logging in, provided you're on the local network. You'll also be able to send and relay mail when connecting from outside the network, providing you set your mail client to use authentication when sending (with the same username and password you use for collecting mail, which is the default with most clients). There's a detailed HOWTO on this at **http://postfix.state-of-mind.de/patrick.koetter/smtpauth**, which also explains how to use TLS to encrypt the traffic. This is important when connecting from outside, otherwise your passwords could be intercepted.

Quick tip

Try to make small changes to configuration files and test them after each change. It makes backtracking after a failure much easier. Making backup copies of the files before each change is a sound insurance policy and *Subversion* is useful for this.

> **"Bayesian spam filters learn what's spam from the mail you receive."**

Spam filtering

Whether you loathe it or merely hate it, spam is something we all have to deal with. Normally, there are two choices: either use your ISP or mail host's spam filters and run the risk of rejecting good mail because of false positives (mails that are incorrectly identified as spam), or use a local spam filter. The latter option gives you more control, but it means you must still download and scan the mails, leading to an annoying delay each time you start up your mail program. Now you have an alternative: let your mail server scan each mail as it comes in by leaving it to run in the background. This means that when you run your mail program, it only has to download the good mails. There are a number of choices here, such as *Spam Assassin* (**http://spamassassin.apache.org**) and *Bogofilter* (**http://bogofilter.sourceforge.net**), but we'll use *Dspam* (**http://dspam.nuclearelephant.com**) here. All are available in most distros' repositories and use a method of finding spam called Bayesian analysis, which involves looking at the occurrence of various words in emails.

There are several ways of calling *Dspam* from *Postfix* – as a content filter, a mailbox transport or a mailbox command. We'll use the last of these, which may not be the most efficient, but it is the most straightforward to set up. You can always look at other methods once you have the spam filtering working as you like it.

Now *Postfix* can call *Dspam* directly for each mail, but this is inefficient, so instead we'll use a *Dspam* daemon that's run as a service when you boot. *Postfix* calls *dspamc*, which uses the daemon instead of running a separate *Dspam* instance. The first step is to edit **/etc/dspam.conf** (it may be in a sub-directory of **/etc**) and ensure the following are set:

```
TrustedDeliveryAgent "/usr/bin/procmail"
UntrustedDeliveryAgent "/usr/bin/procmail -d %u"
Preference "signatureLocation=headers"
ServerMode dspam
ServerPass.Relay1      "secret"
ClientIdent    "secret@Relay1"
```

The first two lines tell *Dspam* to use *Procmail* to deliver the mail for all types of user. The **signatureLocation** option tells *Dspam* to store the mail signature in the message headers, instead of the body. *Dspam* adds a signature to each mail it processes, so it won't try to process the same mail twice. The last three options enable the *dspamc* client to communicate with the daemon. Make sure the *Dspam* daemon is set to start at boot, using your distro's services configuration program, and restart it after editing the configuration file.

Then edit **/etc/postfix/main.cf** and set

```
mailbox_command = /usr/bin/dspamc --client --mode=teft
--deliver=innocent,spam --feature=noise,whitelist --user
$USER -- -d %u
```

which tells *Postfix*, respectively, to use *dspamc* to deliver the mail, to learn from your choices (**mode=teft**), to deliver both non-spam and spam and which user to deliver the mail as. *Dspam* has two main ways of dealing with spam – it can quarantine it in an area that users can access with a web browser to check for false positives and then delete, or it can still deliver the mail but tag it as spam. This is a personal choice, but we prefer the last option, since the mail is still accessible through any mail client and is already sorted into a separate mail folder. You can filter the mail using the

```
X-DSPAM-Result:
```

header, which will contain either **Innocent** or **Spam**.

Restart *Postfix* after editing the configuration file, or force it to reload the configuration by using:

```
postfix reload
```

Bayesian spam filters learn from your email what's spam and what's not. This is what makes them so useful – they don't use a generic set of rules, so they fit in with your needs. The downside, however, is that they need to be taught, which takes place in two ways. First, you have to feed them a lot of

spam and non-spam (occasionally referred to as ham) emails to give them a picture of what you do and don't want, so don't delete any spam you receive before installing *Dspam*, just put it in a separate folder. You can then perform initial training with **dspam_train**, like this:

```
dspam_train username spam_dir ham_dir
```

Because *Dspam* provides individual mail filtering, it maintains a separate classification database for each user. The other two arguments are paths to directories containing the spam and ham mails, in a separate file for each mail. As this is the standard maildir format used by IMAP servers, you can simply pass the paths to the inbox and spam folders of your server, say:

```
dspam_train arthur /var/spool/mail/arthur/.INBOX.spam/cur
/var/spool/mail/arthur/cur
```

You will also need to provide 'on the job' training, particularly in the early days. This involves informing *Dspam* whenever it has incorrectly classified a mail. Being conscientious about this early on will give you much higher success rates in the long run, because *Dspam* assumes that any uncorrected classifications are correct and learns from them. The easiest way to pass false positives or negatives to *Dspam* is directly from your mail client by setting up a couple of mail aliases. These are defined in **/etc/mail/aliases** and this file will already contain some standard mail aliases, so edit it as root and add these two lines to it

```
dspam-spam: "|/usr/bin/dspam --user nobody --source=error
--class=spam"
dspam-notspam: "|/usr/bin/dspam --user nobody
--source=error --class=innocent"
```

An alias definition consists of the alias name, a colon and then the destination. This can be another user, a complete email address or any combination of these separated by commas. It can also be a command preceded by the pipe symbol (**|**), which runs the program and passes it the contents of the mail on standard input. Now you can report false positives or negatives by simply forwarding the mail to **dspam-notspam@your.domain** or **dspam-spam@your.domain** respectively. If your mailer has a macro facility, you can then bind these actions to a hotkey or menu item, which makes training *Dspam* much easier. We use a 'nobody' user to make this a generic alias, but you could also set up a separate pair of aliases for each of your users. After editing the **aliases** file, you'll need to run

```
newaliases
```

to add the changes to *Postfix*'s alias database.

Early spam filtering

Dspam does a good job of detecting spam, but some spam is more easily handled by *Postfix*. Spammers care little for the niceties of email communication, so it's not surprising that they pay scant attention to the normal rules of transmission. In their attempts to send out as many mails as possible, they take shortcuts, omitting or forging information that genuine mails contain. Because of this, you can stop a lot of spam at the initial connection stage of the SMTP conversation by rejecting the connection. This means you save bandwidth, as you don't download the mail and *Dspam* doesn't have to process it. You can reject several of these types of invalid mail by adding this to **/etc/postfix/main.cf**:

```
# HELO restrictions
smtpd_delay_reject = yes
smtpd_helo_required = yes
smtpd_helo_restrictions =
        permit_mynetworks,
        reject_non_fqdn_hostname,
        reject_invalid_hostname,
        permit

# Sender restrictions
smtpd_sender_restrictions =
        permit_sasl_authenticated,
        permit_mynetworks,
        reject_non_fqdn_sender,
        reject_unknown_sender_domain,
        permit

# Recipient restrictions
smtpd_recipient_restrictions =
        reject_unauth_pipelining,
        reject_non_fqdn_recipient,
        reject_unknown_recipient_domain,
        permit_mynetworks,
        permit_sasl_authenticated,
        reject_unauth_destination,
        permit
```

There's no room to go into the detail of how these work here, but if you want to find out more about them, you should head on over to **www.postfix.org**.

Postfix logging

Postfix logs everything to *syslog*, and it can send a lot of output. It's usually best to tell your system logger to send these messages to a separate log file. If you use **syslog-ng** (most distros do), you can add these lines to **/etc/syslog-ng/syslog-ng.conf** to do this.

```
#Postfix logging
destination mail { file("/var/log/mail.log"); };
filter f_mail { facility(mail); };
filter f_notmail { not facility(mail); };
log { source(src); filter(f_mail); destination(mail); };
```

Virus scanning

If Windows users will be downloading mail from your server, it is a good idea to scan incoming mails for viruses as well. *Dspam* can do this for you with the aid of *ClamAV*. First, make sure *ClamAV* is installed and the *clamd* service is running, then edit **dspam.conf** and uncomment the three *ClamAV* settings that follow:

```
ClamAVPort 3310
ClamAVHost 127.0.0.1
ClamAVResponse spam
```

The first two lines should not need changing, while the third tells *Dspam* what to do with the mail when *ClamAV* identifies a virus – **reject** drops the mail and returns an error, **accept** takes the mail and then quietly discards it, and **spam** tells *Dspam* to treat it as spam and quarantine or tag it.

You should now have a complete mail system for SMTP, IMAP and webmail. The programs covered here are immensely flexible and we have covered only the basic options, so roll up your sleeves, read the documentation and tweak the software to suit your needs. LXF

» **Next month** We'll use Virtual Private Networking to extend network reach.

Python: Create

Part 2: Think life could be better? It really could if you follow **Nick Veitch's** advice and create a chatbot to do your bidding.

Our expert

Nick Veitch Nick Veitch launched and edited *Linux Format* for its first eight years. Then he went bad.

Think about how much easier life would be if we had a few servants to do what we bid them to: collect the mail, clean out the gryphon's cage, kick Mike – that sort of thing. Sadly, we're living a little bit too early in the historical timeline for cheap, efficient and compliant robot servants, so we'll have to make do. The next best thing is a virtual servant to do stuff for you. Yes, that could describe practically any piece of software, but what I'm thinking of is something with an easy-to-use interface that can perform a useful task and tell you things you want to know. If it can confound your friends and convince your enemies to reappraise their animosity matrix – well, all the better.

An infrequently used but convenient way of communicating with such an entity would be via chat. After all, why bother with complicated SSH tunnels or boring web-based bidding-doers when you can communicate easily through a medium that you are probably using a lot of the time anyway? In light of that, our little servant will be a chatbot that will sit and wait to hear its master's voice over chat channels, or possibly pop up to let you know about something, should the need arise.

I am Xmpppy

The Jabber/XMPP protocol is well covered by Python. It exists as part of the all-encompassing network module, *Twisted*, but there's also a more lightweight implementation that's suitable for our needs, which goes by the name of *Xmpppy*. You should be able to find a package available for your distro or, alternatively, you can just download the Python code yourself at **http://xmpppy.sourceforge.net**.

To understand how to use *Xmpppy*, we'll start off with a few command line examples, but first you'll need an account to try it out on. Why not register an extra Google account for your bot too? You'll also need at least two Jabber IDs for testing. For our purposes, we created a new account with Gmail for our bot, logged in via a web browser and invited another Gmail account to chat. You can set this up within *Xmpppy* itself, but it becomes complicated – for initial experimentation, it's wise to have two accounts that can already chat to each other. When you're ready, type **python** into the shell to get to the interactive Python command line:

```
>>> import xmpp
>>> jid=xmpp.protocol.JID("botaddress@googlemail.com")
```

OK, the first thing we did is set up an instance of a Jabber ID. In this case it looks pretty much like an email address, because we're using a Google account for this tutorial, but any Jabber ID will do. Effectively there are two parts (there can be three, but we'll discuss that later), which are the

With our guidance, your very own chatbot can soon be offering you personalised abuse.

» **Last month** We mashed up our desktop so it could forecast the weather.

a chatbot

username and the domain. The domain is where *Xmpppy* is going to expect to find the server.

Now we've established the user, we should raise an instance of a client object. In *Xmpppy*, the client is the object that controls the connection, handles the messages and generally interacts with the server. There are a few steps to creating the client and connecting to it. First, we create an instance of the client (which requires a Jabber ID as an argument) and then we try to establish a connection to the server. Once the connection is made, we need to authenticate ourselves before we can actually do anything.

```
>>> myclient=xmpp.Client(jid.getDomain(),debug=[])
>>> myclient.connect()
'tls'
>>> myclient.auth(jid.getNode(),'botpasswd')
'sasl'
```

There are a few things to look at here. To begin with, we've created a client instance and called the **jid.getDomain** method to get the server name from the **jid** object we've already created. You're required to set a debug level as well, which governs the amount of feedback you'll receive. We've set ours to an empty list, but if you really want your screen filled with messages, you could place the string **always** in the list, although this may help if you have problems.

Get connected

After setting our debug level, we created a connection (substitute your own password here) and got the return result **tls**, which means a secure connection has been established with the server – you might also see **tcp** for a standard network connection, or an empty string if the connection fails.

The XMPP/Jabber protocol is in common use, so you can connect to the bot through various clients, *Pidgin* included.

Let's Disco

XMPP includes a framework of plugins called *Disco*. These extend the message protocols to enable the addition of other message types, including SIP (voice), file transfers, and just about anything else you could feasibly imagine a point-to-point network being used for. There's absolutely no reason why our bot couldn't be extended to include this functionality as well, maybe to facilitate easy storage, or to read you the news.

Thankfully, Google has extended its implementation of XMPP with some extra features, which you'll find documented here: **http://code.google.com/apis/talk/jep_extensions/extensions.html**.

Now, depending on your typing speed and your server, the next step may not work. This is because Google's servers expect you to authenticate the connection pretty rapidly, otherwise it will just be dropped, leaving you scratching your head trying to figure out why it hasn't worked. You only have a few seconds of time, so you might want to try joining those lines together from the command line instead:

```
>>> myclient.connect() ; myclient.auth(jid.
getNode(),'botpasswd')
'tls'
'sasl'
```

The **sasl** result tells you that the connection has been authenticated using SASL (Simple Authentication and Security Layer) via the means of your password. There's one final step to initialising the connection, which is to announce your presence. As you know, on Google Chat and other Jabber services, there are a few different levels of client presences (for example, Available or Busy). Apart from anything else, these are used by the server to build rosters of available contacts, and there's a special method in *Xmpppy* to do just that:

```
myclient.sendInitPresence()
```

Note that some servers will forbid you to do anything until you've set this property.

Message in a bottle

Hurrah, we're in! Now, before things get too boring, let's construct and send a chat message to our target (your own Google address could be a worthy candidate). We need to know some things first, though – the address to send it to, and the message itself. Once we have those, we can create a message instance and send it:

```
mymsg=xmpp.protocol.Message("evilbot.1@gmail.
com","hello", "chat")
myclient.send(mymsg)
```

Assuming that you've managed to invite and accept chats from the bot account, you should get a nice surprise.

One thing that might have you scratching your head a bit is how to set your status message, because the server will only display a status state and a message supplied by the »

Quick tip

You can easily tighten your bot's security by only accepting commands from a specific ID.

The *Xmpppy* project page is a bit light on documentation, but the API docs capably explain how the module really works.

client. On Google's server that means you can be not there, Invisible, Away, Do Not Disturb or Available, as well as having a text status message. And that's the key – setting your status is done by sending an XMPP message; a special one that only goes to the server. The message includes both your state and the status text itself. It's a bit puzzling until you know that, but it does make sense – XMPP is designed to transport messages, so why meddle with extra protocols just to send status messages? The server handles updating everyone who is connected to you. As a side note, the XMPP protocol does enable you to send explicit presence messages to individuals but, although there are some exotic uses for this, we're usually happy with the server taking care of

```
>>> presence = xmpp.Presence(status = "Ready!", show =
"chat", priority = '1')
>>> myclient.send(presence)
```

The *Xmpppy* module has a separate way of dealing with presence messages, because they take different arguments, but the mechanism for sending them is exactly the same.

Listening in

To make a proper bot, we also need to get messages coming in. That's a little trickier, to be honest. The *Xmpppy* framework will receive the messages sent to the client object instance you've created and keep them in a stack to be processed. But how do they get processed? Well, *Xmpppy* makes use of a concept called handlers. Before you can process a message, you define a function or method that will act as a receiver for the data. When you're ready to process items in the message stack, you simply call the process method on the client object. This makes it sound more complicated than it is – in other words, you're telling the client where to shove messages, then poking it to put the messages through that function. Of course, we can still do this on the command line in Python, but it gets a little bit messier because we need to write the handler function and then, for normal use, we'd have to create a loop to process messages as they arrive.

We'll demonstrate the code for this on the command line first, but it makes more sense to build it into an object class of its own (as we'll see later on). In this example, we're also going to send ourselves a chat message for testing with the following, but you might want to try it from Gmail instead:

```
>>> def msgparser(instance, message):
...     print "new message!"
...     print "from: " + str(message.getFrom())
...     print "msg: " + str(message.getBody())
...
>>> myclient.RegisterHandler('message', msgparser)
>>> mymsg=xmpp.protocol.Message("evilbot@gmail.
com","hello", "chat")
>>> myclient.send(mymsg)
'5'
>>> myclient.Process(1)
new message!
from: evilbot@gmail.com
msg: hello
1493
>>>
```

As we can see, the parser function is simple. It uses the available **getFrom()** and **getBody()** methods in an incoming

Getting Python help

If you're new to Python, but have experience programming in other languages, there shouldn't be too much to trouble you once you remember to indent your code properly. The main Python site has plenty of documentation explaining the language features and syntax, and also how to use the standard modules that ship with the language.

message, converts them to a string and prints them on the console. For a real bot, we also need to pass the sender into a variable (so we can reply) and perhaps we could choose to further parse the message in order to create a reply.

The voice that commands

For our very simple bot, we will adopt a syntax that means special commands can be passed to the bot by preceding them with '**#**'. Therefore, if a message appears beginning with the hash character, we should try and do something with it. Otherwise, we'll return a random response from a list.

Adding a handler for commands could get large and cumbersome. To save space and effort, the class that we'll define contains a sort of shortcut. Making use of Python's **eval** function, we'll construct a function call to a method in our class with the same name as the command that it's passed in. This is a bit of a fudge, but it saves space and makes it easier to add commands – you just have to define a new method for them. For a more well-rounded chatbot, you would probably want to add a handler mechanism, so that extra commands could be added to instances of the class, but here's how the simpler method would work (note that this code is a class method):

```
def messageparse(self, instance, message):
 m = message.getBody()
 sender=message.getFrom()
 if (m[0]=='#'):
 self.log('command received')
 #special cases to go here
 #general command - expects a method to exist, and to
require the sender id
 try:
 eval('self.'+m[1:]+'(sender)')
 except:
 self.client.send((xmpp.protocol.Message(sender, 'Sorry
Dave, I can\'t do that...')))
 else:
 #i guess we should say something to be polite
 self.client.send((xmpp.protocol.Message(sender, random.
choice(self.responses))))
```

As we see here, the message is checked for **#**, and if there is one, we construct a method call from the rest of the string, include the sender ID and try to execute it. The **try** and **execute** constructions catch exceptions, such as the method being nonexistent. If there's no command, we just send a message back, randomly chosen from a list called **responses**. In the real world, you'll need to import the **random** module for this; **random.choice** just picks a random item from whatever it's given.

An example command handler might then be:

```
def uptime(self, sender):
 import subprocess
 p=subprocess.Popen(["uptime"], stdout=subprocess.PIPE)
 r=p.communicate()[0]
 self.client.send((xmpp.protocol.Message(sender,r)))
```

We should explain this one, because it deals with running commands locally on the machine the bot is running on. Firstly, we have the definition, which accepts its own instance (required in Python), and the sender information the message handler stripped out for us. We have also imported **subprocess** from the standard Python libraries to handle running a command locally.

We're also using something called the Popen method (see the line beginning with **p=subprocess.Popen**), which is thoroughly documented and explained at: **http://docs.python.org/library/subprocess.html**. In brief, we're passing the command to be executed and requesting that the standard output is connected to a pipe. We're then able to connect to the output using the **communicate** method of the Popen instance to get the returned output from the command. The final line packages this up and sends it out as a chat message. So now you can query your server through a chat client and have it abuse you!

Going further

The complete, commented code for our chatbot class is included on the **LXF**DVD. If you run the code directly, make sure you change it to include the login details of your own bot. However, if you import it as a module instead, you can just pass the account and password details through the constructor. The example code should give you enough information to see how it all works.

You will probably want to add your own methods to the class to define further commands. There's no reason why you can't use these commands to either automate some tasks, or query external resources – you could easily create a reminder system, for example, that queries an online calendar. Alternatively, why not plug your chatbot into a translation service, so you can include it in a conversation and have it translate what you say into a different language?

The chatbot is really just a conduit – an easy-to-use interface that enables you to communicate with the script. What the chatbot actually does is up to you and what you program it to do. Please feel free to write in and let us know if you come up with some more uses – and remember to check out the full code on the DVD if you're having trouble understanding something. LXF

Python versions

Python 3.0 is available, but since it's code-incompatible with earlier versions, most distros are (and will be for the foreseeable future) still using 2.x as the default. Hence all our code in this series is, or will be, 2.x compatible to make it easy for the majority of users.

Python editors

Python is particular about syntax. That's no bad thing, but it can be problematic if you're trying to write code with inappropriate editors. We use *Vim* and *Kate*, both of which have syntax highlighting and other features that make it easy to bash out a script in Python that works. You could also use *Eric*, the Python-based IDE, which includes various other Python-specific features. This is available as a package for mainstream distros, or you can get it from **http://eric-ide.python-projects.org/eric4-download.html**.

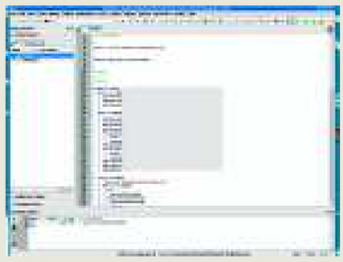

Either use an editor with Python-specific syntax hightlighting, or get the language-specific *Eric* IDE.

Next month Twitter... bringing Python power to 140-character messaging.

Security: Protect

Martin Meredith teaches you how to manage your ports, deal with vulnerabilities and stop hackers from taking advantage of your server.

Our expert

Martin Meredith is a Debian and Ubuntu developer, and a security expert for a major UK online retailer

Way back in the early days of dial-up, the internet mostly contained library catalogues, military secrets and students' Dungeons and Dragons spec sheets. Now there are websites for people, their pets, their friends and family, and their businesses. However, while most people are happy to use a free hosting provider, or to pay a company to host their websites for them, the more dedicated web master tends to plump up for a dedicated server, or a Virtual Private Server (VPS).

Running your own server means that you have to be aware of the multitude of potential security issues you're exposed to on the internet, though. These days, most home computers have a firewall in place, or connect through a router that can protect them from the dangers lurking on the web. If you own a server, you'll still need a firewall, but there's much more you can do to be safe online and we'll show you how.

How secure are you?

On the internet, every service you connect to has a port that it uses. For example, when you connect to a website, you use port 80 (or port 443 for HTTPS) and when you SSH, you use port 22. FTP uses port 21, IMAP uses port 143, and so on. When a server runs, it opens that port and waits for an incoming connection.

So, how does this affect you? When you run a server, you might have a few different services running – maybe you have a basic LAMP stack, or an mail server. These services normally open up their ports for anyone on the internet to see, which isn't always what we want.

By way of example, let's have a look at a recently set up server using a program called *nmap*. This is normally available from your distribution's package manager, or at **http://nmap.org**. Once you have *nmap* installed, supply it with your server's address to get an output that looks like:

```
mez@lazy: % nmap torpor

Starting Nmap 4.76 ( http://nmap.org ) at 2009-05-04 11:56 BST
Interesting ports on torpor:
Not shown: 984 closed ports
PORT      STATE    SERVICE
22/tcp    open     ssh
25/tcp    open     smtp
53/tcp    open     domain
80/tcp    open     http
110/tcp   open     pop3
143/tcp   open     imap
993/tcp   open     imaps
1234/tcp  open     hotline
3306/tcp  open     mysql
10000/tcp open     snet-sensor-mgmt

Nmap done: 1 IP address (1 host up) scanned in 3.46 seconds
```

The results of *nmap* show 10 ports are open for anyone on the internet to connect to. Most of these are normal things that you'd want to see available on the server, such as SSH and email. Some of them, however, we don't want to have open to the internet (such as *MySQL*) and some of them are just plain confusing, such as port 1234.

Is anybody listening?

So, as we've seen above, we can't always tell what's listening in on a certain port. A quick Google search tells us that port 1234 is normally used by various trojan viruses, which means we should probably do a little more investigation to find out what's going on and why that port would need to be open.

The easiest way to find out what program is listening on a specific port is to use the **netstat** command. Running this without any options will give you a list of the currently open connections. We, however, want to find out what is listening to a specific port.

To find out what program is running on port 1234, we'll need to run **netstat -pnl** as root. This will give us a list that might look a little confusing at first, but we're only interested in two of the columns in the output. Since we already know

» **Last month** We built and configured a home telephone network with Asterisk.

your server

the port number we're looking for, all we need to look at is the column called Local Address for an entry ending with **:1234** (the colon separates the IP address and the port number). The line we're interested in is:

```
tcp        0      0 0.0.0.0:1234            0.0.0.0:*               LISTEN
24481/php-cgi
```

This tells us that the program *php-cgi* is listening on all IP addresses on the server (0.0.0.0 means any IP address) for a connection from anywhere. PHP is a scripting language, and here we've set it to listen in Fast CGI mode to port 1234.

Anyone being able to run scripts on our server is not a good thing, so it's time to enlist the help of *iptables*.

Laying down the law

Available by default on nearly all Linux distributions, *iptables* is Linux's answer to all your firewalling needs. In essence, a firewall sits between your computer and the internet, either denying or allowing traffic that's going to and from your server, based on a set of rules.

iptables is relatively easy to set up, but learning how to write the rules can take time. Below is the rules file that we'll be using on this server:

```
*filter
:INPUT DROP [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -m state --state RELATED,ESTABLISHED -j
ACCEPT
-A INPUT -s 127.0.0.1 -d 127.0.0.1 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 22 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 25 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 53 -j ACCEPT
-A INPUT -p udp -m udp --dport 53 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 80 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 110 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 143 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 993 -j ACCEPT
-A INPUT -p tcp -m tcp --dport 10000 -j ACCEPT
-A INPUT -s 192.168.1.3 -p tcp -m tcp –dport 3306 -j
ACCEPT
-A INPUT -p icmp -m icmp --icmp-type 0 -j ACCEPT
-A INPUT -p icmp -m icmp --icmp-type 8 -j ACCEPT
COMMIT
```

Here, ***filter** tells us which table these rules apply to. We'll only be covering the filter table here, but other tables are available for setting up processes such as NAT (Network Address Translation), routing and so on.

The next few lines, those containing **[0:0]**, set up the default policies for the chains in *iptables*. When we use *iptables*, we generally work with three chains: **INPUT** for all incoming connections, **FORWARD** for connections that are forwarded to another server (we won't be using this here), and **OUTPUT** for all outgoing connections. A policy for these chains can be set to **ACCEPT**, **DROP**, or **REJECT**, which either allows the connection, ignores the connection attempt, or sends back an error code saying that the port is not open.

Here we've set the **ACCEPT** and **FORWARD** chains to **DROP** all incoming connections, and the **OUTPUT** chain to **ACCEPT** by default. However, these defaults are only run when a packet doesn't match any of the rules and even then we'll need to add a few exceptions.

Before we do that, it's worth mentioning that the final line, **COMMIT**, tells *iptables* to apply these rules to the firewall. This is the point where your firewall becomes active and starts defending your server.

Plugging the holes

The lines beginning **-A INPUT** set up our firewall rules. These are written as if you were calling *iptables* directly, but without the **iptables** at the start of the command. You can find out more information about how to write these rules from the *iptables* man page.

Each rule starts with **-A INPUT**, which tells *iptables* to append the rule to the chain for input. If we were adding rules to govern output instead, all we'd need to do is substitute **INPUT** for **OUTPUT** here. After that, we have the matching part of the rule and we end the rule line with **-j ACCEPT**, which tells *iptables* to join the packet to the **ACCEPT** chain – in other words, let the packet through. If we wanted to **REJECT** or **DROP** the packet instead, we could substitute those terms in the place of **ACCEPT** accordingly.

Now let's take a look at the line:

```
-A INPUT -m state --state RELATED,ESTABLISHED -j
ACCEPT
```

which tells *iptables* to allow any connection that's related or established to be allowed through. This is a good rule to have, »

Quick tip

Running **netstat -pnl** as root will tell you which programs are listening to ports.

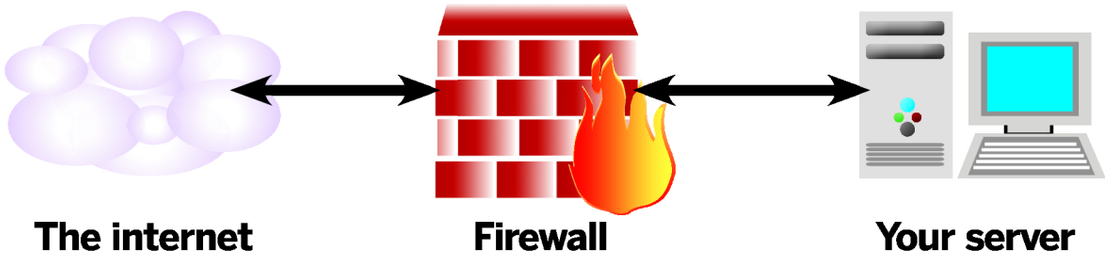

You've seen this diagram before I'll wager, and yes, it really is this simple.

since it permits any connection you have already made to continue – for example, the current SSH connection will stay active if you're using SSH to change the firewall rules. It also allows incoming connections that are related to outgoing connections (TCP works both ways and you'll find you can't make outgoing connections from the box without this rule).

The next line down tells *iptables* to accept any connections that originate from 127.0.0.1 (**-s** means source IP address) connecting to 127.0.0.1 (**-d** refers to the destination IP address), which translates to 'accept any connection I make to myself'.

In the next nine lines, we open up the ports that we saw earlier to the outside world. Notice that we don't have any code that allows access to ports 3306 (*MySQL*) and 1234 (the misconfigured PHP installation), because the firewall rule that allows connections locally takes these into account. However, we do want to add an exception to rejecting all outside access to *MySQL*, which comes in the form of:

```
-A INPUT -s 192.168.1.3 -p tcp -m tcp –dport 3306 -j ACCEPT
```

This specifically enables the IP address 192.168.1.3 to connect to port 3306 and it's here because we have another server on our local network that needs to access *MySQL*.

The final two lines above **COMMIT** allow the server to respond to pings, and enable it to receive ping responses.

After saving the file (to a location similar to **/etc/iptables.conf**) we can run the command:

```
iptables-restore < /etc/iptables.conf
```

enabling the firewall rules we've written above. If we run *nmap* again, it gives us the following results

```
mez@lazy: % sudo nmap torpor
Starting Nmap 4.76 ( http://nmap.org ) at 2009-05-04 11:56 BST
Interesting ports on torpor:
Not shown: 984 filtered ports
PORT      STATE    SERVICE
22/tcp    open     ssh
25/tcp    open     smtp
53/tcp    open     domain
80/tcp    open     http
110/tcp   open     pop3
143/tcp   open     imap
993/tcp   open     imaps
10000/tcp open     snet-sensor-mgmt
Nmap done: 1 IP address (1 host up) scanned in 10.30 seconds
```

Notice that the output now says that it's not showing filtered ports instead of closed ports. Also, you may have to run *nmap* as root – in our tests, the above firewall configuration worked too well and *nmap* couldn't find all the open ports without being run as root.

Quick tip

You can make your firewall rules run on startup by adding the **iptables-restore** command to **/etc/rc.local**.

Et tu, Brutus?

Even with different security measures in place, there's still the possibility that someone can compromise your server simply by guessing the password! There are a number of scripts out there that use a dictionary to try every possible combination of username and password in a relatively short time.

Fail2ban (**www.fail2ban.org**) tries to overcome this by monitoring your log files for failed login attempts. When it detects that someone has tried to login multiple times and failed, it will restrict their access to the server (using the firewall) for a certain amount of time.

This isn't a great solution if you regularly forget your password, but it is good for stopping people in their tracks when they try to use brute force to pry their way into your server.

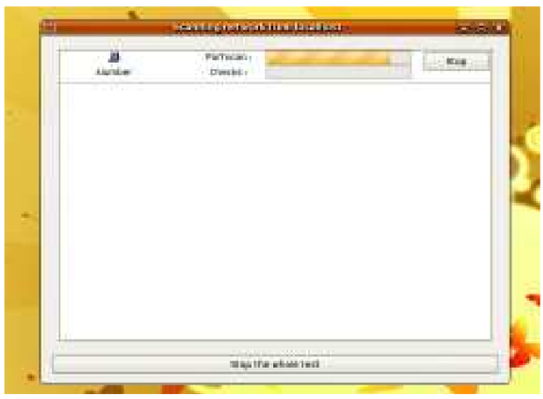

› ***Nessus* scans your servers for vulnerabilities in its code that could be exploited by the unscrupulous.**

What's Nessus-ary?

This basic firewall should protect us to some extent, but there are other potential vulnerabilities in every system, especially for servers that don't get updated often. To combat this, there is a tool known as *Nessus*, which comes in two parts – the client and the server. The reason for this is so that the server can be installed in a remote location and used to test the connection to the local service. In this case, the computer that we're testing isn't local to us, so we can install both the server and the client locally.

Once you've installed *Nessus* and *nessusd* from your preferred package manager, you'll need to set up a user. Run the command **nessus-adduser** and follow the instructions on screen. Now, run the *Nessus* client (found under Applications > Internet > Nessus in Ubuntu).

When *Nessus* runs, it works through a list of plugins that enable the testing of various vulnerabilities. Because there are new vulnerabilities discovered every day, it's wise to keep *Nessus* up to date. Tenable Network Security, the application's creator, has two different lists of plugins available – the HomeFeed and the ProfessionalFeed. The ProfessionalFeed provides a more current list, meaning that the newest vulnerabilities can be checked as soon as plugins for them exist. ProfessionalFeed costs $1,200, though, so unless you're working with *Nessus* on a daily basis, or working in a high-security environment, HomeFeed is easily enough.

Tenable requires you to register for its feed downloads, which you can do at **http://linkpot.net/enviably**. After that, run the command **nessus-fetch –register <your registration code>** with the registration code that you'll be supplied via email.

Once that's set up and you load the program, you may find the *Nessus*'s main screen is a bit overwhelming. Instead of trying to walk you through all of it abstractly, we'll show you how to perform a basic scan.

First of all, you need to log in to the server with the credentials that you set up previously. After you've logged in, you'll be presented with the Plugins screen. This is a list of scripts that *Nessus* will try to run. Click the Select All button, and flip to the Scan Options tab. Here, you can find the general scan options for the *Nessus* scan. The most

interesting of these options is the Safe Checks checkbox, because *Nessus* won't perform scans that have the potential to crash your server if it's ticked. However, as long as you have physical access to the server, or some way of resetting it remotely, we'd recommend leaving this box empty, because this enables you to perform a more in-depth scan. Ultimately, we'd rather have *Nessus* find a security hole and crash the server than leave the vulnerability in place.

Finally, all we need to do is put the name or IP address of a server into the Target tab and then getting *Nessus* underway is a simple case of clicking Start The Scan and letting the program do the rest.

Once *Nessus* has finished scanning, it will present you with a report showing a list of items that *Nessus* has found. A lot of the things that *Nessus* returns are informational messages such as 'you have an SSH server running', but anything important will have a red stop icon next to it. If *Nessus* finds vulnerabilities, it will tell you how to fix them (or point you to the right place to find out how to fix them). You should act on any critical points as soon as possible.

Trip 'em up

If all goes well with your *Nessus* scan and subsequent vulnerability fixing, you should now have a server that's hard for someone to get into. However, that doesn't mean that it's impossible to breach your security – the only way to make sure that nobody can get into your system it to turn it off.

So, what can we do if we can't guarantee the security of a system unless it's off, or at least disconnected from the internet? At the point where a hacker gets into your system, you can still, at a minimum, make sure you know that they've been there. An Intrusion Detection System (IDS) enables you to do this, although we hope you'll never have to rely on it.

The most widely known IDS at the moment is *Tripwire*, which has been around since 1992. At the moment, three programs called *Tripwire* are available and they all do the same thing. Only one of them, however, is open source. You can find it in your distribution's package manager, or by heading over to **http://tripwire.sf.net**.

Tripwire works by creating a database of all the files on your system and notifying you when something about them has changed. Because of the way the system works, the best time to install *Tripwire* is before you connect the system to the internet. If someone has worked their way into your system before the software has been installed, then *Tripwire* will just end up making sure that any back door the attacker has created stays in place.

Setting up *Tripwire* is a lengthy process, and somewhat beyond the scope of this tutorial, but you can find more information about the process at **www.alwanza.com/howto/linux/tripwire.html** and **www.tripwire.com**.

Keep an eye on your servers

Possibly the most important thing to do when considering security for your servers is to make sure that you keep an eye on what's going on. This process is generally known as auditing and, if done correctly, can ensure that issues are resolved before they have the opportunity to become blown out of proportion.

While we've already used *Nessus* to audit the security from outside the server and mentioned using *Tripwire* to audit the integrity of files on the system, there are many other auditing tools available. However, we've found that the best way of auditing a system is to read the log files – or least get a program to do that for you (more on that in a moment).

Servers tend to generate a lot of information about what's happening. If you have a look in your **/var/log** folder, you'll find a variety of different log files, ranging from the system log to *Apache* access logs. These logs provide plenty of useful information, but figuring out what's worth reading and what's not can be a tedious job.

This is where *Logwatch* comes in – it's a utility that reads your log files, and can send you daily emails about the most interesting parts. These emails may be long, but they will let you know when things aren't going to plan. They can also keep you informed if one of your users starts trying to access things that they shouldn't be, giving you the time to address the issue before they cause a problem.

Under Debian or Ubuntu, the installation and configuration of *Logwatch* is simple. However, to receive the information as emails in HTML format, you'll need to alter the output method, output format, and the email address you're sending the information to. You can change these options by editing **/usr/share/logwatch/default.conf/logwatch.conf** so that:

```
Output = mail
Format = html
MailTo = youraddress@yourdomain.net
```

This will send the reports to your email inbox, meaning you can stay up to date with server news wherever you are. LXF

A command line alternative

In this tutorial, we've spent some time teaching you how to write firewall rules on the command line. There are, however, alternatives to this approach. Our favourite (and one that we tend to use on a regular basis) is *Webmin*, a tool that enables you to manage your server through your web browser. If you've installed *Webmin* on your server, you can find the firewall options under Networking > Firewall.

If you've understood the basics of creating rules for a firewall that we've covered here, using *Webmin* should be a walk in the park. To begin with, click Revert Configuration to load the firewall rules that you set up earlier. Then you can modify the rules as appropriate for your new *Webmin* setup. Finally, you should hit Apply Configuration to finalise the modified rules once you've finished changing them.

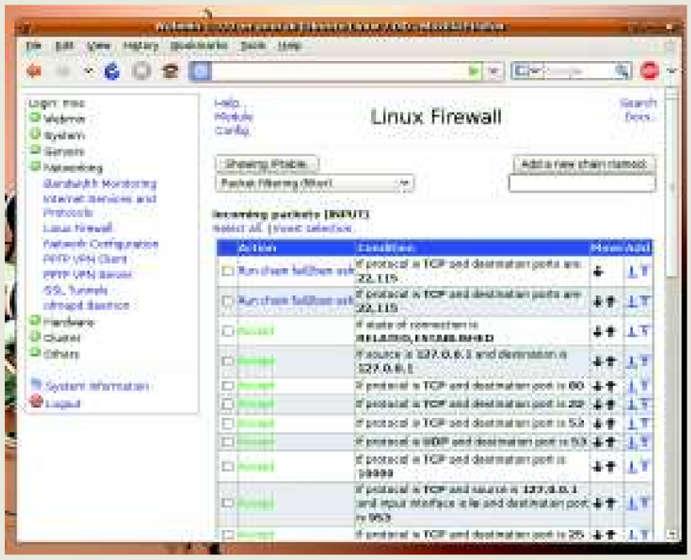

» **Next month** Get total privacy on the web with the Incognito live CD.

Subscribe today

And save up to 30%!

PLUS Get a set of fantastic Linux Format stickers!

LAST CHANCE TO GET THESE EXCLUSIVE STICKERS!

UK Pay only £14.75 By Direct Debit every 3 months!

Step by step: Using your LXF stickers

1 Tainted
Like so many others, Mike's Linux box was marred with the mark of Microsoft...

2 Painted
... until we crafted the finest and most Linux-friendly stickers known to man!

3 Sainted
With the *Linux Format* stickers in place, Mike's code is faster, better and bug-free. Huzzah!

OVERSEAS PRICES	
	1 year (13 issues)
EU	£100.80
Rest of the world	£108

www.myfavouritemagazines.co.uk/lxd/P074

Please use the full web address shown to enjoy this special offer

techradar.com
Deep into technology

NEWS | REVIEWS | BLOGS | FORUMS

See all news | **Browse news categories**

Where am I? News ▸ News by technology ▸ World of tech

WORLD OF TECH

21 guilty pleasures for geeks

Things you don't even admit on anonymous forums

It's OK, you're in safe company. We know you didn't really buy that gadget because it could interface seamlessly with the cloud, or because it offers the best range of features for Small to Medium Enterprises, or even because it has the fastest processor this side of a supercomputer. You just like touching its rubbery buttons and the fact that it goes 'ping'.

You're not alone. In fact, here are some other guilty pleasures that we'd like to share with you...

1. The iPhone Unlock swipe

You've used your iPhone every day since you bought it, but even all those apps you've downloaded can't compare to the first thing you do with your iPhone every day: gently, caressingly, swiping your finger from left to right across the Unlock button. It's the nearest you ever get to foreplay. And you love it.

2. Seeing how juiced up your PC is in Vista / Windows 7

You know you've got a kicking system, and how. Yet every day you dive into the Welcome Centre just to see it all by the numbers. "Windows-capable? Don't make us laugh. Just look at the stats I'm getting," you say (probably to yourself).

3. The Macintosh bong

A cross-between a call to action, geek whimsy and the harbinger of PC doom, the Macintosh bong is one of the great geek pleasures. It confirms your status as a man (or woman) who has thunk different, sticking it to the (PC) man with your non-conformist ways. Sometimes you stick it through your kick-ass speakers. Sometimes you simply soak it up on headphones. But the best is when you Zap the PRAM by holding down the Command-Option-P-R keys so your Mac starts up again, and again.... and again. Or you could just watch the Bong video on YouTube.

4. Looking at your blog's user stats

OK, so your mum is the only person who looks at your blog and that's only because you set her homepage to it last time your configured her PC, but still. Every day you pile on to Statcounter, Technorati or Google Analytics just to see how well you're doing. What, you're supposed to get them to ignore your own IP address? Where's the pleasure in that?

5. Waggling the analogue sticks on your PS3

Whether you're fragging your way through FPS mayhem, careering around a virtual racetrack or lolloping around a platformer, the best bit about it all is the feel of those rubbery pointy things that squirm beneath your thumbs. You especially love it when you waggle so far that the joystick reaches its softly clunking stop. It's even better when you send it yawing back the other way.

6. Cranking up your home cinema system for the THX OMMMMMmmmmm

200W per channel amplifier? Check. Seven surround sound speakers and two active subs? Check. The neighbours are out? Stuff it, who cares! That latest Blu-ray blockbuster may shift more CGI pixels than a supercomputer in a Pixar render-farm, but hey – just listen to this. Everyone, come on, OMMMMMmmmmm. It's what war would sound like if Hare Krishna had an army.

7. Smashing the ball in Wii Tennis

OK, so you could just wave your arm around lazily while lying semi-comatose on your sofa. But surely it's much (much) better to lift the Wii-mote right above your head, violently jerk your arm down and then listen out for that thundering whack as your virtual racket hits the virtual ball. You're rewarded with the shimmering, heady sight of the ball screaming banshee-like across the net. Bingo, one set-winning ace in the bag and the shame-faced joy of seeing your four-year-old nemesis crumple, crying, to the floor. Guilty pleasure? Oh yes.

8. Deliberately trying to confuse your sat nav

You've plotted the fastest route from here to Maplin, only now you're just going to take a teensy detour. "Perform a legal U-turn", TessTess says politely. "Perform a legal U-turn", she says again, tension rising slightly. "I SAID, PERFORM A LEGAL U-TURN! What's the matter, can't you tell your left from right, or something?" TessTess, of course, has failed to spot that you're driving a DeLorean. And this is now 1955...

There's more! Read the full article at www.techradar.com/565083

WHAT IS TECHRADAR?

- Latest tech news
- Expert reviews
- Tutorials and tips

ALSO ON TECHRADAR

- 50 seriously useful Windows 7 tips
- The PC of 2059
- 14 best Linux distros
- Best free iPhone apps
- Future tech
- Asus Eee PC 1000HE reviewed
- 7 Sky innovations coming to your television

See techradar.com for more info...

Come on, admit it: you love swiping your iPhone screen

Zoom

Answers

Got a question about open source? Whatever your level, email it to **lxf.answers@futurenet.co.uk** for a solution.

This month we answer questions on:

1 Tardy file previews
2 SSD filesystems
3 Remote backups
4 Tips for documentation
5 USB modems
6 X Windows problems
7 awk and sed
8 Retrieving a missing MBR
★ Automatic CD duplication

1 Empty home

Q **After upgrading Ubuntu to 8.04 with kernel linux-2.6.24-23-generic and Gnome 2.22.3, I have experienced a few small problems on my Dell Inspiron 6400 laptop. The Gnome file browser gets stuck when I try to display my home directory. It works fine on all other directories on the filesystem.**

***Audacity* doesn't let me play recordings. I have tried every possible configuration of output device. When performing the upgrade to 8.04 I got a warning about the package manager being corrupt – something to do with not being able to configure *hplip*.**

Jon Spragg

A How long have you left the file browser to display the home directory? If you have file previews enabled it may be taking its time generating previews for one or more files, especially if any of them are on a network share and not the local machine. To test this, either turn off previews completely in the File Browser preferences by setting each file type to Never, or at least make sure they are set to Local Files Only with a small maximum size.

If this makes a difference, create a temporary directory and move files and directories from your home directory into it until the delays stop. Then you'll know which file or files are causing the problem and be able to take appropriate action.

Can you play audio files from other programs, or save an audio file from *Audacity* and then play it from another program? You need to try this to determine whether the problem lies with *Audacity* or your audio playback in general. Whatever program fails to play the audio, try running it from a terminal. That won't make it play the file, but you will be able to read any error messages from it.

The error from the package manager probably came from a corrupt package, or one with an invalid signature. Refreshing the package list in *Synaptic* and installing any updated packages should clear this.

2 Eee disk life

Q **I own an Eee PC 900 on which I have recently installed Ubuntu Netbook Remix 9.04. Besides now having Ubuntu's flexibility and updates at my fingertips, I am also in love with the NBR interface, which works surprisingly well after some minor tweaks (including a patched kernel). However, I have read some worrying posts about NBR (and other alternate distros) not being very easy on the Eee PCs built-in solid state disk. People suggest all kinds of precautions to be taken when installing such a distro, including:**

1. **Never choose to use a journalling file system on the SSD partitions.**
2. **Never use a swap partition on the SSD.**
3. **Always edit your new installation fstab to mount the SSD partitions with *noatime*.**
4. **Never log messages or errors to the SSD.**

There are other suggestions too, concerning the behaviour of certain applications, like *Firefox*'s cache. I can't remember reading about this SSD-problem in your tutorial on replacing the operating system on the Eee with something different (LXF109). So, can you tell me if any (or all) of these suggestion make sense? What seems weird to me is that a default install of NBR on the 900 gives me an ext3 filesystem, a swap partition and mounts with *relatime*. Don't they care about SSD lifespan at Canonical, or is the whole issue just rubbish?

Marc Floris

› It can take a while to generate some file previews, so stick to smaller, local files.

A These fears are all based on the fact that the SSD is essentially flash memory, which has a limited write cycle. But it is not used in the same was as flash memory is used in memory cards and USB sticks. The problem with flash memory is that each cell can only handle a limited number of write operations, USB devices and memory cards are normally rated for between 100,000 »

Our experts

» Whatever your question, we can find an expert to answer it. From installation and modem woes to network administrations, we will get the answer for you – just fire off a letter or email and it'll all be taken care of.

Neil Bothwick
Having run a small ISP and produced coverdiscs for *Linux Format*, Neil describes himself as a Linux jack of all trades.

Mike Saunders
Mike was one of the original contributors to LXF prototype *Linux Answers*. His specialisms are programming, window managers, *init* scripts and the SNES.

Paul Hudson
Paul is the in-house programming whizz, but is also ready and willing to tackle all your web- and database-related problems.

Graham Morrison
When not reviewing mountains of software and tinkering with *MythTV*, Graham is on call to answer any hardware and virtualisation issues!

WHERE TO SEND YOUR QUESTIONS:

Write to *Linux Format*, Future Publishing, 30 Monmouth Street, Bath BA1 2BW or send an email to lxf.answers@futurenet.co.uk.

LINUX FORMAT
The #1 source for Linux

Back issues

Missed one?

Every month *Linux Format* brings you the best tutorials, the essential reviews and the latest news. But if you've missed a must-read feature or a vital program from our expertly compiled DVDs, order your back issue NOW! And remember, you'll never miss an issue of *Linux Format* if you subscribe (see pages 76 and 102 for details).

Issue 120
July 2009
Product code:
LXFDB0120

In the magazine
Jaunty Jackalope arrives, so we take it for a spin and speak to Mark Shuttleworth about the future of Ubuntu. Also, we help you get organised with *Basket Note Pads* and *Picard*, find out what Intel's doing for Linux, and learn how to code web mashups with the ever-popular Python.

LXFDVD highlights
Itching to try Ubuntu 9.04? Well, we've got an exclusive version just for you. There's also Mandriva 2009 Spring, Apress PDFs and much more.

Issue 119
June 2009
Product code:
LXFDB0119

In the magazine
Kick out the productivity jams and get to work, with our 26 top tips for getting more done with Linux. Plus there's a reminder of why Slackware is so good, the sinister side of JavaScript, we hear from behind the scenes at PHP and we get all old-school with retro file managers.

LXFDVD highlights
Newbies can wet their whistle with PCLinuxOS 2009.1 and Mepis 8, while there's Zenwalk 6 and Slackware 12.2 for the more advanced crowd.

Issue 118
May 2009
Product code:
LXFDB0118

In the magazine
Whatever kind of user you are, there's a Linux distribution to suit – but how do you pick from 323 different options? Well, that's where our cover feature comes in, providing all you need to choose the right distro. If that's too heavy for you, we also find plenty of fun in the amorphous form of *World of Goo*.

LXFDVD highlights
Debian 5 arrived full of great software, delivered with a neat graphical installer and positively oozing power. *Xfce 4.6.0* and tons of apps are also on board.

Issue 117
April 2009
Product code:
LXFDB0117

In the magazine
As overwhelmingly fantastic as it is, Linux can and does go wrong. Fear not – we have the solutions to all your problems! Plus we explore the best hardware to take Linux into your lounge, look at the benefits of *GnuPG* and investigate Moonlight – our answer to MS Silverlight.

LXFDVD highlights
Knoppix 6, the awesome live DVD that allows you to take Linux everywhere, 15 different distros to try out, coding tutorials and 40 apps.

Issue 116
March 2009
Product code:
LXFDB0116

In the magazine
You've sent us letters, we've read them – the result is eight pages of KDE hacks to make your desktop run like a dream. Plus we kick off a new tutorial series dedicated to the Acer Aspire One, show you how to bring data back from the dead and safeguard your data DVDs from scratches.

LXFDVD highlights
OpenSUSE 11.1, the latest SUSE community offering; Slackware 12.2 for you old-school Linux users and a fresh Ubuntu respin with Linux Mint 6.

Issue 115
February 2009
Product code:
LXFDB0115

In the magazine
Virtualisation – it's not just for übergeeks. Anyone can do it, as long as they read our cover feature first. Also in the mag there's a slightly sinister webcams tutorial, our essential security checklist and a guide to speeding up your machine with one of Linux's lighter desktops.

LXFDVD highlights
All hail the Ubuntu-beating Fedora 10: it's got a great community, thousands of packages and bags of shiny new eye candy.

To order, visit www.myfavouritemagazines.co.uk

Select Linux Format from the 'Select your favourite magazine' drop-down box on the right.

Or call the back issues hotline on 0870 837 4773 or +44 1858 438795 for overseas orders

Quote the issue code shown above and have your credit or debit card details ready

and a million writes, depending on the quality of the individual device. This sounds a lot, but some parts of a disk are written to very frequently, like FAT tables and filesystem journals. A damaged journal can be worked around, but a corrupt file allocation table on a FAT filesystem is close to terminal, and this is the filesystem used by removable flash devices.

SSDs are different for a number of reasons. They generally use higher-quality components, netbooks don't use FAT filesystems and, most importantly, SSDs incorporate wear levelling. This means that the load is spread across the 'disk' and writes are not repeatedly made to the same sector.

I have been using journalled filesystems (ext3 and xfs) and swap on my Eee PC900 for over a year. It is used every day and runs Gentoo testing so packages are updated almost daily. Combine that with extensive email and web usage (mailer caches are written to as much as browser caches) and the only disk errors I've had were on the SD card, the only component that doesn't use wear levelling (and was of unknown make and quality). You will need a swap partition (or suitable file) if you want to hibernate your laptop.

Bear in mind that the Eee comes with a two-year warranty, and Asus isn't likely to include technology that is likely to fail in that timeframe. Canonical isn't in the business of breaking hardware (nor is any other distro maker), although I would question the use of *atime* when mounting a filesystem, but that is for performance reasons more than reliability – I use *noatime* on hard disks too. Apart from that, I would be happy to use the NBR setup.

3 Remote backup server

Q I am looking at setting up a remote backup server to back up various Windows desktops and servers. I am looking for a backup package for the Linux Server as well as the client software to manage the backup from the client PCs.
George Lianos

A I suggest you take a look at *BackupPC* (**http://backuppc.sourceforge.net**). It is an entirely server-based backup program. By that I mean that no special software is required on the client PCs and all backups are controlled from and initiated by the server. This means that you do not need to rely on users remembering to start backups, or set up *Cron* tasks on each computer.

BackupPC has a web-based front-end, so you can browse and restore complete backups or individual files or directories. Restoration can be directly back to the filesystem from which the files came, or the files to be restored can be downloaded as a tarball or zip archive. The web interface provides an overview of all the computers it manages, with details of all their backups. You even get email alerts if any backup fails, although *BackupPC* will normally continue when it can with no need for intervention (this occasionally happens to me when I shut down my laptop while it is being backed up).

› *BackupPC* handles backups for many machines, without the need for any special software on those computers.

BackupPC communicates with the client computers using *Samba*, SSH, NFS or *rsync*, so you don't need any special software on the client computers – just ensure that *BackupPC* has permission to read shares or connect via SSH. As you are backing up a number of possibly similar computers, you'll be pleased to know that *BackupPC* saves on both disk space and backup time by storing multiple copies of files as hard links. If the same file is on 10 computers, the server stores only one copy.

There isn't the space to go into a detailed explanation of how to set up *BackupPC* here, but its documentation is extensive and we will be covering setting up a backup server with this in our networking tutorial series in a couple of months.

4 Help with documentation

Q In your April 2009 issue [LXF117] you give a very concise solution to the ".dmrc file is being ignored" permissions problem in the Answers section. I'm suffering from this at present and found your two lines quick and simple.

```
chown -R user: ~user
chmod -R u+rw,go-w ~user
```

I'm not a programmer, but would like to get involved in improving usability by changing the current error message to include this probable solution. How would I do this?
John Stiles

A Improving documentation is one area where users can always help improve open source software, even if they have no programming experience. There are two main ways of offering help and improvements. The less formal method is to post on the software's public mailing list or web forum (almost every project has one) with your suggestions. This may or may not elicit a response from the developers, but it could »

› See page 108 for Rackspace's star letter.

result in a discussion with other users. The other option is to file a bug on the project's bug tracker, which most also have. Despite the name, bug trackers are not used exclusively for bugs; enhancements and feature requests can be submitted too. As all requests are tracked, you and the developers can see when any action is taken on your request.

In a case like this, a bug report with a suitable replacement message is in order, but keep the message short. When a longer message is needed, it may be better to put that elsewhere and have the error message refer to it for more information. If you can provide a patch to the project's source code, so much the better. For documentation bugs, this generally involves downloading the latest source tarball and using *grep* to find the source file containing the message you wish to change. Making a patch file is a straightforward task. Let's say you want to change the file **errors.h** in the current directory. Make a copy called **error.h.orig** then make your changes to **errors.h**. Now create a patch with

```
diff -u errors.h.orig errors.h >errors.h.patch
```

If you read the patch file, you'll see it is a set of instructions to the *patch* command on which lines to remove and add from the old file to turn it into the new one. Submit the patch file with your bug report so the developers can reproduce your changes. Do not submit the complete new file, as they may have made other changes. A patch file lets them merge your changes into the latest version.

5 USB if you must

Q I am using a computer with an AMD 1.6GHz processor. I have installed Ubuntu 8.10 in a separate partition from Windows XP. My internet provider is BT and the USB modem is a BT Voyager 105. This is where my problem starts, as BT tell me they don't recognise Ubuntu and cannot suggest how I can access the net. I downloaded some gobbledygook from the Ubuntu forums but that did not work. Surely there is a simpler, step-by-step way to get the internet in Ubuntu.

By the way, my Lexmark printer would not work so I had to splash out on an HP 4100 printer, so don't ask me to buy too much software as I am a pensioner and at present skint!

Graham Phillips-Lewis

Setting up a Voyager USB modem is a bit of a fiddle, but is worthwhile if you cannot get an Ethernet ADSL modem.

A I know you don't want to spend any money, but a few pounds spent on a decent modem will save you so much trouble. The best thing you can say about the free modems given away by ISPs is that they are reasonable value for money! A decent modem connects by Ethernet, not USB, and needs no special drivers or software on the computer. The standard networking stack and a web browser, which everything has by default, will do just fine. Most standalone modems also include a router and firewall, making your system more secure whichever operating system it uses. Because a proper modem handles the network protocols internally instead of offloading the work to the computer's CPU with a driver, both the network connection and your computer in general are more responsive. Such a modem can be bought for around £20.

If you stick with the Voyager USB modem, you have to accept that it will work less efficiently and that you will have some work to get it set up. This is true on Windows too, but you do get automated driver installation there. To do this with Ubuntu 8.10, you need to download two files. As your connection is not working under Linux, do this in Windows or on a different computer. Go to **http://eciadsl.flashtux.org/download.php** and fetch the Ubuntu package, currently **eciadsl-usermode_0.12-1_i386.deb**. Then go to **http://archive.ubuntu.com/ubuntu/pool/universe/r/rp-pppoe** and get the latest i386 Deb file, which is currently **pppoe_3.8-3_i386.deb**. The version number in these files may change if updates were released after we wrote this. Copy these files to a USB stick and transfer it to your Ubuntu computer. After making sure the modem isn't plugged in, install each of the packages by double-clicking them, the **pppoe** file first.

Now you have to configure the modem with the settings for your ISP (this is the part that Windows users get done for them). In your terminal, run

```
sudo eciadsl-config-tk
```

to open the graphical configuration program (if it fails to run use **eciadsl-config-text**). At the top of the window, set your username and

Star Question Winner!

This month's winner is dwilmot20. Your new Kodak camera is on its way!

CD duplication

Q I'm trying to build a CD duplicator from an old Sempron-based machine with four IDE CD-RW drives and a SATA hard disk. I've installed Ubuntu 9.04 and tried *GnomeBaker* and *K3b* but they don't seem to support multiple CD burning. We need to produce about 200 CDs for various projects at the school where I work and I wondered if you had any ideas that a newbie could cope with.

dwilmot20, from the forums

A I'd use a shell script for this, after creating the ISO image with whichever program you prefer. A script is far better suited to such a repetitive task than having to keep pressing GUI buttons. You could use something like this

```
#!/bin/sh
DEVICES="/dev/cdrom0 /dev/cdrom1 /dev/cdrom2 /dev/cdrom3"
for DEV in DEVICES; do
  cdrecord -eject dev=$DEV "$1" &
done
```

Type this into a text editor, such as *Gedit*, list your CD writer devices in the **DEVICES** line, save the script somewhere in your path – say, **/usr/local/bin/multiburn.sh** or **~/bin/multiburn.sh –** and make it executable. If you save it in the **bin** folder in your home directory, you can use your file manager to set the permissions: right-click on the file and select Properties. Otherwise, set the permissions in a terminal with

```
sudo chmod +x /usr/local/bin/multiburn.sh
```

Assuming you have made the ISO image with *K3b* or whichever mastering program you prefer, put a blank CD in each of the drives and run the script, giving it the path to the ISO image.

```
multiburn.sh /path/to/image.iso
```

When all four discs have ejected, replace them and run the command again. You could modify the script to run again after a keypress. There are only two keypresses to rerun it (Up and Enter) so it's not a time saver, but it may be a good learning exercise if you're so minded.

» password to those given to you by BT and set VPI and VCI to 0 and 38 respectively. Pick the correct modem from the list, set the PPP mode to VCM_RFC2364, click on the Remove Dabusb button (you can safely ignore any messages it gives you) followed by Create Config. This should open a window containing a number of messages ending in OK. Your modem is now installed and set up, and you can plug in the modem and connect to the internet by running

```
eciadsl-start
```

You can attach this command to a desktop icon by right-clicking on the desktop and selecting Create Launcher. Put **eciadsl-start** (or **eciadsl-stop**) in both the name and command boxes. You can also have the command run automatically when the desktop opens by going to System > Preferences > Sessions, clicking on the Add button and entering the program name.

6 Portability or not

Q Whatever happened to the concept of portability in the Linux community? I know that different distributions will work on a variety of unique machines. Well, shouldn't the same concept be applied to the distribution itself when considering upgrades? I have an Acer Aspire 64-bit dual processor machine with a GeForce 8200 Graphics card, and everything works in Fedora 9. I must admit that I have to poke around with the screen resolution a bit to get what I want, but that's minor. When I try to upgrade to Fedora 10 64-bit, *X Window System* fails to start – there is no GUI, just a login prompt, and startx doesn't do anything. Why would *X* work in Fedora 9 but not Fedora 10? Shouldn't the same stuff work in the same operating system even if it's an upgrade?

Willim Brown

A The answer to your question is "yes" but I suspect you were hoping for a little more. I doubt that **startx** does nothing; it may not start *X*, but it will report to the terminal what has gone wrong. You can find further information in the log file at **/var/log/Xorg.0.log**. Errors are marked with "(EE)", and you can extract them from the general information in this file with

```
grep EE /var/log/Xorg.0.log
```

To some extent, distros are at the mercy of the various software developers. In this case, the changes to *X.org* are the most likely culprit. The *X.org* team are trying to move away from the cumbersome and sometimes cryptic **xorg.conf** file to a system based entirely on auto-detection, the idea being that any hardware you have should just work, which yours clearly does not. The quickest solution, if you still have your old Fedora 9 installation, or a backup of the important files, is to copy **/etc/X11/xorg.conf** from Fedora 9 to your Fedora 10 setup. *X.org* respects anything in this file in preference to auto-detected settings, so it should start working again.

If you don't have an **xorg.conf** file, you can create one for your hardware by running the following as root:

```
yum install system-config-display
system-config-display
```

This presents the same display setup GUI as you saw during installation of Fedora 9, and creates **/etc/X11/xorg.conf** based on your choices. If it still fails, because it cannot detect suitable settings for your hardware for example, you can specify your choices on the command line, for example

```
system-config-display --set-resolution=1024x768 --set-depth=24 --set-driver=nvidia
```

Run **system-config-display --help** to see all the options.

7 Sedding easy

Q I'm writing a *Bash* script that, among other things, installs assorted packages. I'm using *Zenity* to make it look sexy in a GUI sort of way and running it on 64-bit Debian 5.0.

Unfortunately, if one of the packages I want to install is available on the Lenny installation DVD, the thing just sits there: under the hood, it has said 'please insert DVD', but *Zenity* doesn't see that and doesn't relay that. So to the user, it just looks as if the thing has hung. The relevant bit of code is as follows:

```
Select all
  ( echo "33"
```

Frequently asked questions...

Remote desktops

» I want to connect to my home computer from work and run a graphical program, like my email client, is this possible?
Yes it is. There are two basic ways of doing this, with several variants on each.

» What are the different options?
The first is *X* forwarding, which is often best when running Linux on both computers. This works over an SSH connection, where you run a program in the remote shell and it is displayed on your local desktop. For example

```
ssh -X me@my.home.computer
kmail
```

The **-X** tells SSH to use *X* forwarding.

» Is that all I need to do to run the program on my local computer?
Strictly speaking, the program is running on the remote computer, but it is displaying its window on the local system. You may need to edit **/etc/ssh/sshd_config** to set 'X11Forwarding' to 'yes'.

» What if I want to display the whole desktop – can *X* forwarding to that?
Yes it can, by opening a new desktop session that displays on your local computer.

» That's not the same, is it? If I left a program running and want to get at its window, this won't let me, will it?
No, you need a remote desktop program for that, as you would if your local computer were running Windows. The most common option here is VNC (Virtual Network Computer). This is a server/client system, where you have a VNC server running on your remote computer, and run the client on your local system. There are packages at **www.realvnc.com** for all popular OSes, so you can access your Linux desktop from a Windows or Mac system, and vice versa.

» What software will I need to do that?
Which desktop do you use? If you run KDE, you already have all you need as KDE has built-in RFB (remote frame buffer) software. RFB is the protocol used by VNC and compatible systems. With any other desktop you'll have to install the VNC or *TightVNC* package from your distro's repositories.

» Is bandwidth important? Can I use this over a modem link?
Displaying a graphical desktop and keeping up to date with changes at the remote end requires a lot of bandwidth, a modem link is not suitable. A broadband connection is feasible, but remember that no matter how fast your downstream ADSL speed, upstream is unlikely to be more than 384kb/s, so TightVNC is the better option here. Reducing the screen size and depth, getting rid of fancy effects and images and anything else that reduces the amount of data to be transferred will help.

› Check it out: here's a KDE desktop, running on Windows XP via VNC.

```
    aptitude -y install gcc
    echo "66"
    aptitude -y install sysstat
    echo "99"
    ) | zenity --auto-close --progress
--text="Fetching software..." --title="Installing
Software " --width 300
```

A number of questions arise. First, is there a way to make *Zenity* 'event aware' so that it displays messages like 'please insert your installation DVD'? I don't think there is, but I could well be wrong. If so that would be the ideal solution and you don't have to read this question to the end (just let me know the relevant piece of magic!).

Alternatively, I could write code that says 'if they've got the installation DVD in their /etc/ apt/sources.list, comment it out. And 'if they haven't got http://volatile.debian.org/debian- volatile in their sources, add it'. Unfortunately, to analyse the textual contents of files, I can feel heavy doses of *awk* and *sed* coming on, neither of which make me feel particularly clever, inspired or confident! In fact, I'm clueless about both, despite reading the man pages and Google articles about them until I'm cross-eyed.

Would you therefore be able to suggest some *Bash* code that tests for the existence of the two sources in the sources.list file, removes the DVD one if present, adds the online one if it's missing and doesn't add it if it is already present?!

HJR, from the forums

Is the "please insert" message being sent to standard error rather than standard output? If so, you should be able to capture it by adding **2>&1** to your **aptitude** call. Then you should look at the output for the relevant string before passing it to *Zenity*.

You are quite correct in your assumption that modifying text files from the command line needs *sed* or *awk* – *sed* in this case – but these are tools well worth learning. They can be daunting at first, but are indispensable once you get the hang of them. To comment out any CD/DVD source from **sources.list**, you would use this command

```
sed  -i 's/^deb cdrom/# deb cdrom/' /etc/apt/
sources.list
```

To break this down into manageable chunks, the **-i** means 'replace the existing file with the modified contents'; the **s** means 'substitute anything matching the first string with the second'. So this replaces any instance of **deb cdrom** at the start of a line with **# deb cdrom**. This is the same way that *Synaptic* modifies the file so the source can be re-enabled in there should the need arise.

Adding a source line is as simple as echoing it to the file, but messing with someone's **sources** file is not particularly friendly, so you could check whether you need to make the modifications with *grep*, back up the file before changing it and restore the backup when you've finished.

```
cp /etc/apt/sources.list /etc/apt/sources.list.$$
sed  -i 's/^deb cdrom/# deb cdrom/' /etc/apt/
sources.list
grep -q "^deb http://volatile.debian.org/debian-
volatile" /etc/apt/sources.list || echo
"repository line" >>/etc/apt/sources.list
# do your stuff here
if diff -q /etc/apt/sources.list /etc/apt/sources.
list.$$
then
  rm /etc/apt/sources.list
else
  mv /etc/apt/sources.list.$$ /etc/apt/sources.
list
fi
```

$$ is the current process number, giving a unique(ish) name for the backup file. The **grep** command means your repository is only added if missing, and the final part checks whether **sources.list** has changed, and restores the old one if so. »

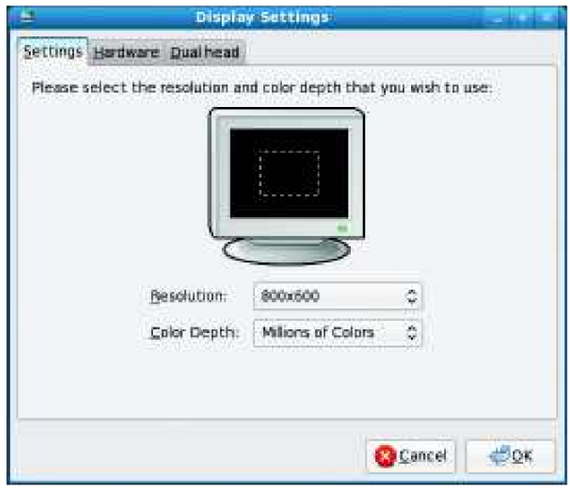

› Fedora's display configuration program will run without *X*, which is useful, as you most need it is when you don't have *X*.

A quick reference to...

Initrd

If you look in your /boot directory or your bootloader menu file, you will see references to initrd files. These are ramdisk images (some are ramfs images, which is a newer version of the same principle). A ramdisk, as the name implies, is a disk-like storage device that lives entirely in memory, and the initrd file is that ramdisk saved to a file. That's what it is, but what is it for?

The initrd, or initial ramdisk, is a ramdisk that is loaded by the kernel when it starts up. This ramdisk becomes the root filesystem, and scripts are run from here to set up the system before passing control to the real root partition on your hard disk. The main function of a ramdisk is to load kernel modules. Distros are built to work on the largest possible range of hardware, which means that most driver modules are built. If these were compiled into the kernel image, you'd get a huge kernel that was slow to load, took up too much memory and was 90% redundant – that's why Linux has loadable kernel modules. You can't load a modules needed to mount the root filesystem from the root filesystem, so Linux needs a way to load the drivers for the hard disk controller and filesystem, and maybe some other bits like *LVM* or *dm-crypt*, beforehand. This is what the ramdisk does, the **linuxrc** script on this disk loads the relevant modules, runs any setup programs that may be needed (such as for *LVM* or an encrypted root filesystem) and then switches to the hard disk root.

An initrd image is a filesystem in a cpio file that has been compressed with *gzip*, so you can unpack, mount, modify and repack one with these lines:

```
cd /mnt/tmp
zcat /path/to/initrd | cpio -id
#modify files here
find . -depth | cpio -o | gzip >/path/to/
newinitrd
```

Help us to help you

In order to give the most complete answers to your questions, we need to know as much as possible. If you get an error message, please tell us the exact message and precisely what you did to invoke it. If you have a hardware problem, let us know about the hardware. If Linux is already running, you can use the excellent *Hardinfo* program (**http://hardinfo.berlios.de**) that gives a full report on your hardware and system. The report is an HTML file that you can attach to your mail.

If you're unwilling, or unable, to install *Hardinfo*, run the following commands in a root terminal and attach the **system.txt** file to your email. This will still be a great help in diagnosing your problem.

```
uname -a >system.txt
lspci >>system.txt
lspci -vv >>system.txt
```

The Big Question How do I get my Ubuntu installation back?

8 Bootloader recovery

Q I recently installed Ubuntu 9.04, but made the mistake of installing XP in another partition, and had forgotten that XP tramps over the MBR. This meant that I could no longer access the email files on my Ubuntu installation.

I thought then that if I were to install another copy of Ubuntu this might give me access to the MBR and perhaps give me a way of getting to the first Ubuntu installation via a modification to the MBR.

Can I access the email folders in the first installation using my second installation, or can you advise how I can regain access the first Ubuntu installation?
Patricia

A Only the initial bootloader code is stored in the MBR; the files needed to present the boot menu and boot your choice are all stored in the boot partition of your original Ubuntu install. You can restore the *Grub* MBR from most live CDs, and some are designed for this task, like SuperGrub Disc (**www.supergrubdisk.org**) or System Rescue CD (**www.sysresccd.org**). Super Grub Disc can automatically fix the MBR in the case on one Windows and one Linux installation on the computer. Even with two Linux installations, it's only a case of following the menus and choosing the distro to use as the bootloader config. There are detailed instructions at **www.supergrubdisk.org/wiki/Howto_Fix_Grub**.

You can fix *Grub* from any live CD with a few commands. First you need to determine which disk contains your boot directory, *Grub* counts drives and partitions from zero, so the second partition on the first drive – **/dev/sda2** in Linux terms – is called (hd0,1). Once you know that, open a terminal and run

```
sudo grub
root (hd0,1)
setup (hd0)
quit
```

The first command enters the *Grub* shell as root. The next identifies the partition containing your boot files, then you write the bootloader to the MBR with **setup** and exit the *Grub* shell. If you're unsure of which partition to use, the *Grub* shell has a **find** command to help. Most distros create a symlink to the kernel from **/boot/vmlinuz**, so after entering the *Grub* shell, run

```
find /boot/vmlinuz
```

This will return one result for each of your Ubuntu installations, so pick the right one.

If you want to access the old installation from the new one, you need to make sure the user and group IDs match. While you may have used the same username on both installations, the filesystem stores only the numeric ID, so if your user had UID 1000 on the first install and 1001 on the second, you won't have full read/write access to the other installation's files. To make them readable on the current installation, run this command

```
chown -R username: /other/home/username
```

replacing the username and mount point with the appropriate values.

Step by step: Restore a broken MBR

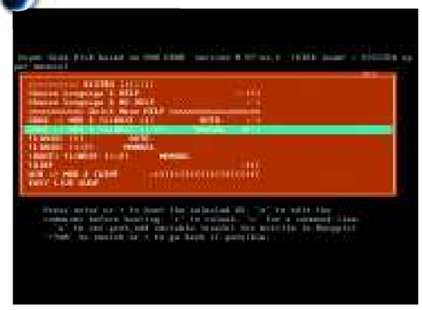

1 Super Grub Disk

Installed to a CD or USB stick, this boots straight into a menu that gives various options for fixing your MBR and bootloader.

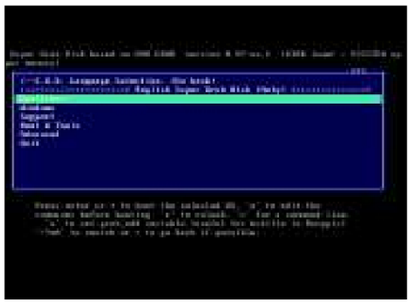

2 Online help

Before diving in to hack your MBR, take a few moments to read the documentation for the options you are considering.

3 Options explained

The help entry for each option explains what it does; press any key to go straight to that option.

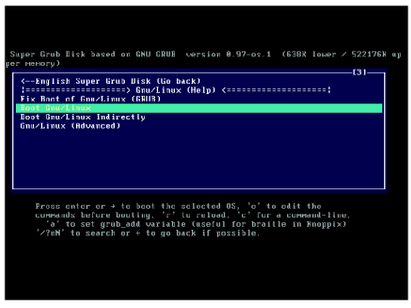

4 Boot Linux

If you don't want to change the MBR now, you can use Super Grub Disk to boot into any of the operating systems on your disk.

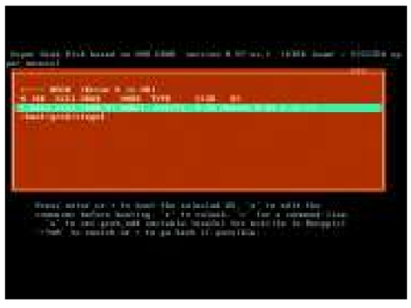

5 Fix the MBR

If you have more than one Linux installation you need to pick the one that holds the bootloader configuration files.

6 Check permissions

If you want to access two installations using the same login details, use the *chown* command to change their numerical IDs. LXF

LINUX FORMAT

The #1 source for Linux

Future Publishing Ltd, 30 Monmouth Street, Bath BA1 2BW **Tel** 01225 442244
Email linuxformat@futurenet.com

ABC **28,734** January–December 2008
A member of the Audited Bureau of Circulations.

EDITORIAL

Editor Paul Hudson
paul.hudson@futurenet.com
Reviews editor Graham Morrison
graham.morrison@futurenet.com
New media editor Mike Saunders
mike.saunders@futurenet.com
Art editor Efrain Hernandez-Mendoza
efrain.hernandez-mendoza @futurenet.com
Operations editor Andrew Gregory
andrew.gregory@futurenet.com
Editorial contributors Jeremy Allison, Neil Bothwick, Chris Brown, David Cartwright, Andy Channelle, Matthew Clapham, Marco Fioretti, Michael J Hammel, Andy Hudson, Juliet Kemp, Susan Linton, Michael Meeks, Martin Meredith, Simon Pickstock, Mayank Sharma, Shashank Sharma, Nick Veitch.
Art assistance Stacey Black, Carlton Hibbert.
Illustrations Shane Collinge, Chris Hedley, Chris Winn.

ADVERTISING SALES

Advertising sales director
Rob Elms 020 7042 4272
rob.elms@futurenet.com
Advertising sales manager
Charlie Said 020 7042 4142
charlie.said@futurenet.com
Account manager
Robert Williams 020 7042 4251
robert.williams@futurenet.com
Sales executive
Katherine Fort 020 7042 4152
katherine.fort@futurenet.com

MANAGERIAL & MARKETING

Brand manager Jennifer Wagner
Promotion executive Emily Bonney
Campaign manager Lee Peacock
Publishing director Stuart Anderton
UK chief operating officer Simon Wear
UK chief executive Robert Price

PRODUCTION AND DISTRIBUTION

Production coordinator Craig Broadbridge
Trade marketing manager Stuart Brown
Distributed by Seymour Distribution Ltd, 2 East Poultry Avenue, London EC1A 9PT **Tel** 020 7429 4000
Overseas distribution by Future Publishing Ltd
Tel +44 (0)1225 442244

LICENSING

Head of international licensing Tim Hudson
thudson@futurenet.co.uk **Tel** +44 (0)1225 788253

CIRCULATION

International account manager John Lawton
john.lawton@futurenet.com

SUBSCRIPTIONS & BACK ISSUES

UK readers' order line 0844 848 2852
General UK enquiries 0844 848 1602
Overseas readers +44 (0) 1858 438795
Online enquiries www.myfavouritemagazines.co.uk
Email *linuxformatdvd1@subscription.co.uk*

Printed in the UK by Midway Clark (Holt) and William Gibbons

Future MEDIA WITH PASSION

Future Publishing Ltd is part of Future plc.
Future produces carefully targeted special-interest magazines for people who share a passion. We aim to satisfy that passion by creating titles offering value for money, reliable information, smart buying advice and which are a pleasure to read. Today we publish more than 150 magazines in the UK, US, France and Italy. Over 100 international editions of our magazines are also published in 30 other countries across the world. Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR).

Non-executive chairman Roger Parry **Chief executive** Stevie Spring
Group finance director John Bowman **Tel** +44 (0)1225 442244 **www.futureplc.com**

ATLANTA • BATH • LONDON • NEW YORK • PARIS • SAN DIEGO • SAN FRANCISCO

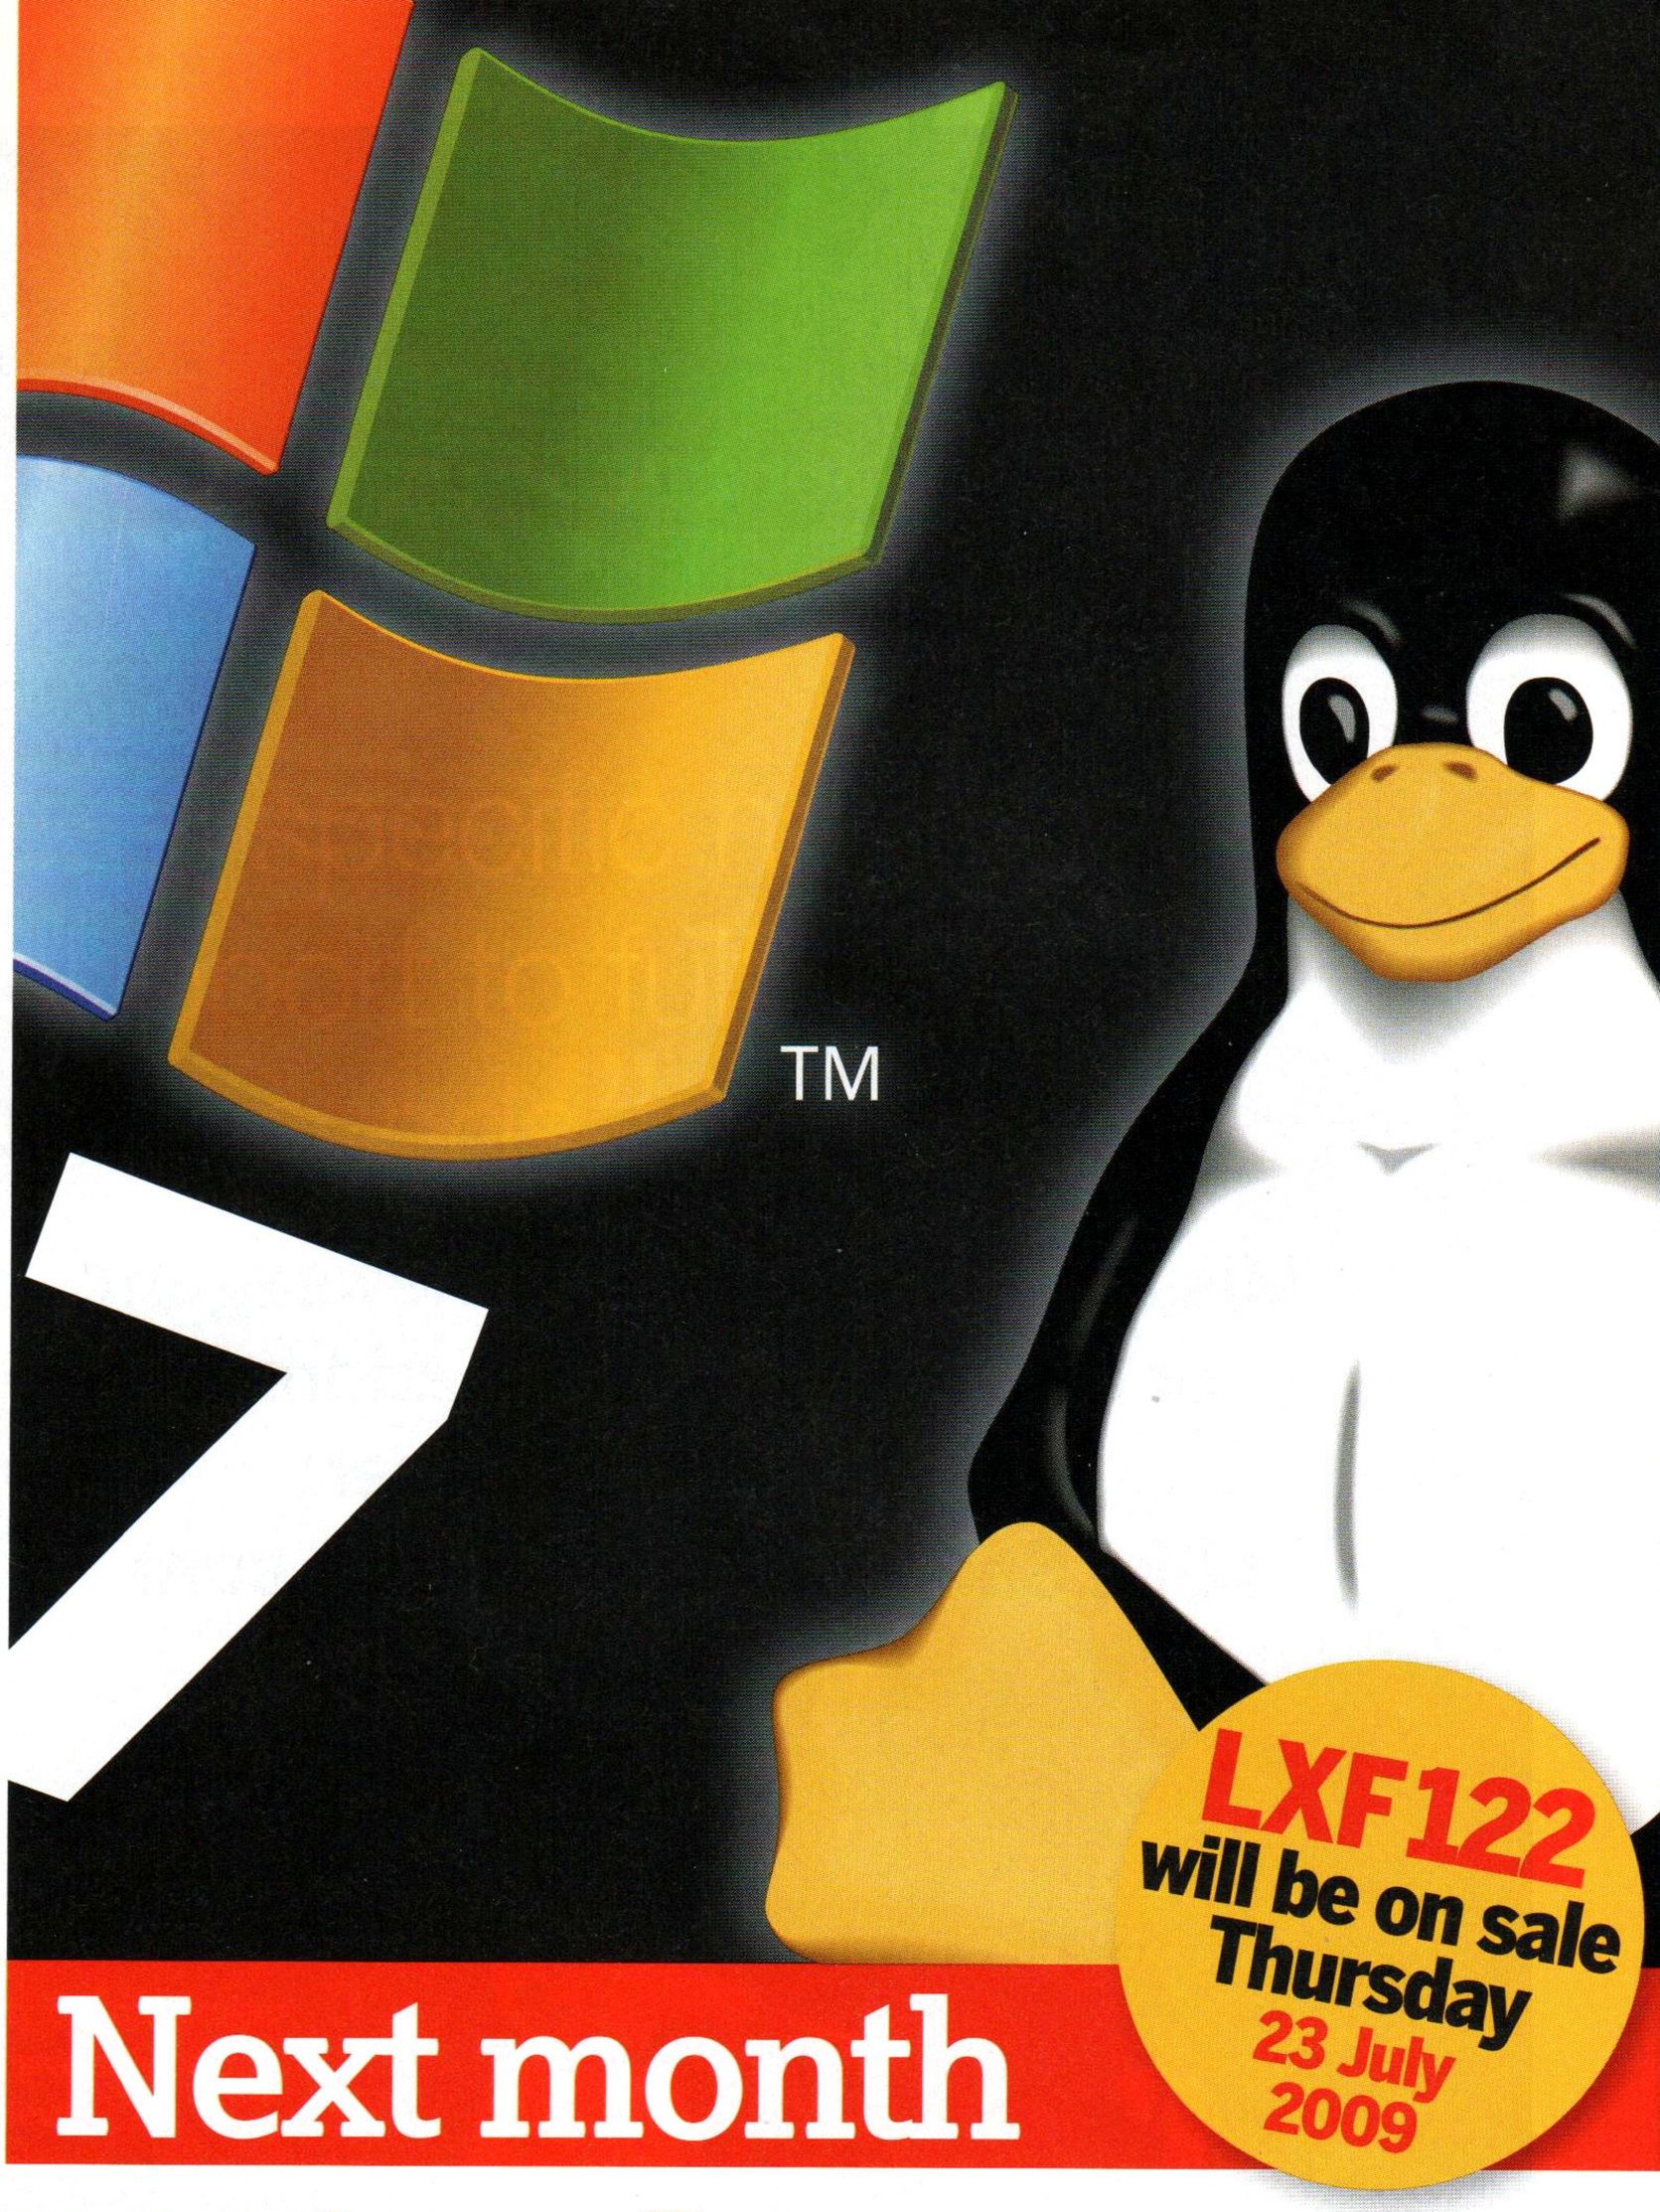

Next month

Windows 7 vs Linux

Beat your friends to new features with free software, without compromising on low system requirements.

Online anonymity explained

Incognito means no one can watch your movements online no matter how hard they try – learn how!

KOffice 2.0

After years of work, the latest KDE office suite is finally here – but has it done enough to compete with *OOo*?

Contents of future issues are subject to change – freedom of choice is of course an advantage of Linux!

LINUX FORMAT
The #1 source for Linux

August 2009

DVD121

Open index.html for full DVD listings

Missing disc? Ask your newsagent or email us with your address and the issue number:
Mike.Saunders@futurenet.com

NEW!

EcoDisc DVD ROM

» Uses fewer materials

» Takes less energy to manufacture

» Tougher and more flexible

WARNING
Do not use this disc in Apple slot drives
Contact support@futurenet.co.uk if you can't run this disc

» This month's DVD contents

DISTROS
Linux Mint 7
OpenSolaris 2009.06
Ulteo OVD

DESKTOP
Alarm Clock 1.0
Audacity 1.2.6
EnergyXT demo
Kdenlive 0.7.4
Me TV 0.9.0
OpenOffice.org 3.1.0
Paperbox 0.4.3
Sox 14.2.0
Sweep 0.9.3
Viewnior 0.4

GAMES
Max Reloaded 0.2.5
Quadra 1.2.0

Stendhal 0.74
Tubularix 0.1.5.2
SuperTuxKart 0.6.1a
Ubuntu Tweak 0.4.7.1

DEVELOPMENT
Eric4 4.3.4
Gaphas 0.4.0
GiftWrap 0.1

INTERNET
BaShare 0.4.2
Pidgin 2.5.6
Twitim 1.2

HOTPICKS
Atomic Worm 0.19
BoPlanets 1.3
Cactus Jukebox 0.4.1
Flush 0.5.1
Geeqie 1.0beta1
Gentoo 0.15.4
Gnome Schedule 2.1.0
Minimum Profit 5.1.2

SERVER
Dspam 3.8.0
Postfix 2.6.1
Sockso 1.1.8
Tripwire 2.4.1.2

ESSENTIALS
ATI driver 9.4
HardInfo 0.5
Kernel 2.6.29.4
Nvidia driver 180.60
Smart Boot Manager 3.7

SYSTEM
VirtualBox 2.2.4
Wireshark 1.0.8
YADSync 1.0

MAGAZINE
Tutorial code
TuxRadar podcasts

HELP
Linux Answers
Newbie guides

LINUX FORMAT
The #1 source for Linux

OpenOffice.org 3.1
GUI updates aplenty

OpenSolaris
Broaden your Unix skills with version 2009.06

MORE!
» Me TV
» Tubularix
» Newbie guides
» VirtualBox 2.2.4

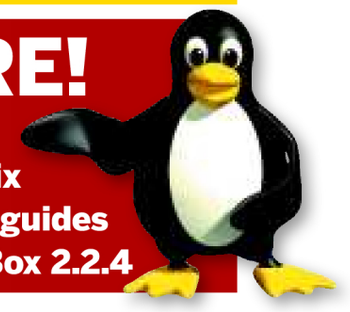

HOSTWAY

"We need highly specific hardware and flexible support to fulfil our own clients' needs"

Hostway don't believe Enterprise level hosting needs to be complicated. We enable companies that depend on web-based applications and services to focus on what's most important to them: their clients.

Design UK, a leading digital agency, chose an Enterprise Managed Hosting solution from Hostway which allows its clients to quickly roll out online marketing campaigns, optimise their websites, and drive web traffic with minumum hassle and optimum reliability. A load-balanced VMware platform, allied to Pillar SAN storage, ensures their operation never fails them or their clients. And all of this is backed-up 24*7 by Hostway's Enterprise support team.

Whatever the size, duration or complexity of your project, Hostway supply a highly scalable, responsive service that is truly focused on what you want to achieve.
If it's important to you, it's important to us!

at Hostway it's personal

Call us today and talk to one of our experts

0808 180 1880

or visit us online at **www.hostway.co.uk**

Leigh Whitney
Design UK

Read how Hostway helped me at
www.hostway.co.uk/itspersonal

WHO'S BEHIND YOUR WEBSITE?

Domain Registration • Email Hosting • Web Hosting • Ecommerce • VPS • Dedicated Servers • Enterprise Managed Hosting

Call us free on 0808 180 1880 • sales@hostway.co.uk • www.hostway.co.uk

Hostway are proud partners with:

redhat. Fedora PLESK VIRTUOZZO HEWLETT PACKARD Cisco Systems pillar vmware

HOSTWAY
THE HOSTING COMPANY

FANATICAL SUPPORT® MAKES ME FEEL:

"LIKE I'M IN SAFE HANDS"

Michael Howarth, PMC Telcom.

If you've never experienced Rackspace's Fanatical Support and hosting services, our customers just might make you jealous. They say no other hosting supplier makes them feel like we do: Secure. Special. Loved, even.

Fortunately, they like to share their feelings. That's how they've made us the most recommended hosting provider in Europe. For everything from simple web servers to complex online applications and email solutions, talk to the most financially stable and leading hosting provider in the world.

Experience it for yourself at **www.rackspace.co.uk/safe** or by calling us on **0800 840 1279**.

experience ***fanatical support***®

rackspace® HOSTING

CMYK

Linux Mint 7

Universal Edition

LINUX FORMAT DVD121

4GB software-packed DVD

WARNING: Before you even put the DVD in your drive, please make sure you read, understand and agree to the terms and conditions on page 68 in the magazine.

Issue 121 August 2009

LXF121D/08/09 Made in the EU

EcoDisc DVD ROM

LXF121D/08/09

Future MEDIA WITH PASSION

WARNING

Do not use this disc in Apple slot drives

Contact support@futurenet.co.uk if you can't run this disc

This month's DVD highlights

LINUX FORMAT

DISTROS
Linux Mint 7
OpenSolaris 2009.06
Ulteo OVD

DESKTOP
Alarm Clock 1.0
Audacity 1.2.6
EnergyXT demo
Kdenlive 0.7.4
Me TV 0.9.0
OpenOffice.org 3.1.0
Paperbox 0.4.3
Sox 14.2.0
Sweep 0.9.3
Viewnior 0.4

DEVELOPMENT
Eric4 4.3.4
Gaphas 0.4.0
GiftWrap 0.1

GAMES
Max Reloaded 0.2.5
Quadra 1.2.0
Stendhal 0.74
Tubularix 0.1.5.2

HOTPICKS
Atomic Worm 0.19
BoPlanets 1.3
Cactus Jukebox 0.4.1
Flush 0.5.1
Geeqie 1.0beta1
Gentoo 0.15.4
Gnome Schedule 2.1.0
Minimum Profit 5.1.2
SuperTuxKart 0.6.1a
Ubuntu Tweak 0.4.7.1

INTERNET
BaShare 0.4.2
Pidgin 2.5.6
Twitim 1.2

SYSTEM
VirtualBox 2.2.4
Wireshark 1.0.8
YADSync 1.0

SERVER
Dspam 3.8.0
Postfix 2.6.1
Sockso 1.1.8
Tripwire 2.4.1.2

MAGAZINE
Tutorial code
TuxRadar podcasts

HELP
Linux Answers
Newbie guides

WARNING
Do not use this disc in Apple slot drives
Contact support@futurenet.co.uk if you can't run this disc

EcoDisc DVD ROM
LXF121D/08/09

Future MEDIA WITH PASSION

recycle
Respect the Environment
Throwing away old or surplus discs?
If you would like your CDs and DVDs to be recycled rather than put into landfill, please post them to:
Polymer-Reprocessors,
Reeds Lane, Moreton, Wirral, CH46 1DW
Please respect the environment and dispose of waste plastic responsibly.

NEED HELP?

If you have a question about the disc, please phone +44 (0)1225 822743 or visit: **www.linuxformat.com/dvdsupport**

● **Please read the DVD pages in the magazine first!**

This is a Future Technology DVD-ROM Future Publishing Ltd cannot accept responsibility for any disruption, damage and/or loss to your data or computer system that may occur while using this DVD, the programs and data on it. Please note that we can only provide basic advice on using the interface and installing software from the disc. Unfortunately we can't give in-depth help on the applications, or on your hardware or operating system.

Defective discs In the unlikely event of your disc being defective please visit **www.linuxformat.com/dvdsupport** for assistance. If you would prefer to talk to a member of our reader support team please call +44 (0)1225 822743.

4GB

Linux Mint 7

Universal Edition

Ultra-slick Live distro – Ubuntu power with a fresh new interface

Ulteo

Try the user-friendly virtual desktop today!

www.ulteo.com

PLUS: OpenSolaris and OpenOffice.org 3.1

And much more – see index.html for full software descriptions

LXFDVD121 August 2009

4GB

Linux Mint 7

Universal Edition

Ultra-slick Live distro – Ubuntu power with a fresh new interface

LXFDVD121

August 2009

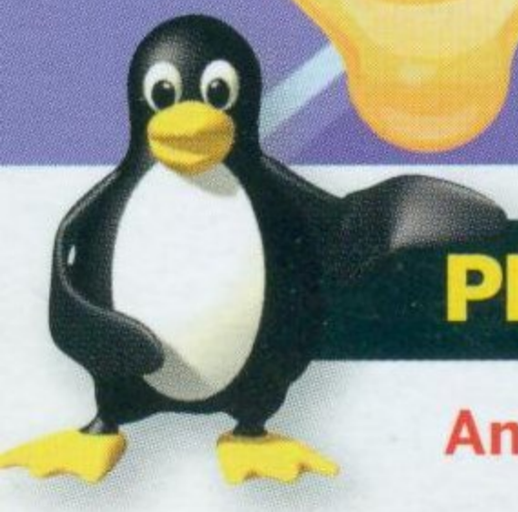

PLUS: OpenSolaris and OpenOffice.org 3.1

And much more – see index.html for full software descriptions

This month's DVD highlights

LINUX FORMAT — The #1 source for Linux

DISTROS
Linux Mint 7
OpenSolaris 2009.06
Ulteo OVD

DESKTOP
Alarm Clock 1.0
Audacity 1.2.6
EnergyXT demo
Kdenlive 0.7.4
Me TV 0.9.0
OpenOffice.org 3.1.0
Paperbox 0.4.3
Sox 14.2.0
Sweep 0.9.3
Viewnior 0.4

DEVELOPMENT
Eric4 4.3.4
Gaphas 0.4.0
GiftWrap 0.1

GAMES
Max Reloaded 0.2.5
Quadra 1.2.0
Stendhal 0.74
Tubularix 0.1.5.2

HOTPICKS
Atomic Worm 0.19
BoPlanets 1.3
Cactus Jukebox 0.4.1
Flush 0.5.1
Geeqie 1.0beta1
Gentoo 0.15.4
Gnome Schedule 2.1.0
Minimum Profit 5.1.2
SuperTuxKart 0.6.1a
Ubuntu Tweak 0.4.7.1

INTERNET
BaShare 0.4.2
Pidgin 2.5.6
Twitim 1.2

SYSTEM
VirtualBox 2.2.4
Wireshark 1.0.8
YADSync 1.0

SERVER
Dspam 3.8.0
Postfix 2.6.1
Sockso 1.1.8
Tripwire 2.4.1.2

MAGAZINE
Tutorial code
TuxRadar podcasts

HELP
Linux Answers
Newbie guides

WARNING
Do not use this disc in Apple slot drives
Contact support@futurenet.co.uk if you can't run this disc

EcoDisc DVD ROM

LXF121D/08/09

Future MEDIA WITH PASSION

recycle

Respect the Environment

Throwing away old or surplus discs?
If you would like your CDs and DVDs to be recycled rather than put into landfill, please post them to:
Polymer-Reprocessors,
Reeds Lane, Moreton, Wirral, CH46 1DW
Please respect the environment and dispose of waste plastic responsibly.

NEED HELP?

If you have a question about the disc, please phone +44 (0)1225 822743 or visit: **www.linuxformat.com/dvdsupport**

● **Please read the DVD pages in the magazine first!**

This is a Future Technology DVD-ROM Future Publishing Ltd cannot accept responsibility for any disruption, damage and/or loss to your data or computer system that may occur while using this DVD, the programs and data on it. Please note that we can only provide basic advice on using the interface and installing software from the disc. Unfortunately we can't give in-depth help on the applications, or on your hardware or operating system.

Defective discs In the unlikely event of your disc being defective please visit **www.linuxformat.com/dvdsupport** for assistance. If you would prefer to talk to a member of our reader support team please call +44 (0)1225 822743.

Linux Mint 7

Universal Edition

LINUX FORMAT The #1 source for Linux

DVD121

4GB software-packed DVD

WARNING: Before you even put the DVD in your drive, please make sure you read, understand and agree to the terms and conditions on page 68 in the magazine.

Issue 121 August 2009

LXF121D/08/09 Made in the EU

EcoDisc
DVD ROM

LXF121D/08/09

Future
MEDIA WITH PASSION

WARNING
Do not use this disc in Apple slot drives

Contact support@futurenet.co.uk if you can't run this disc

Rigel

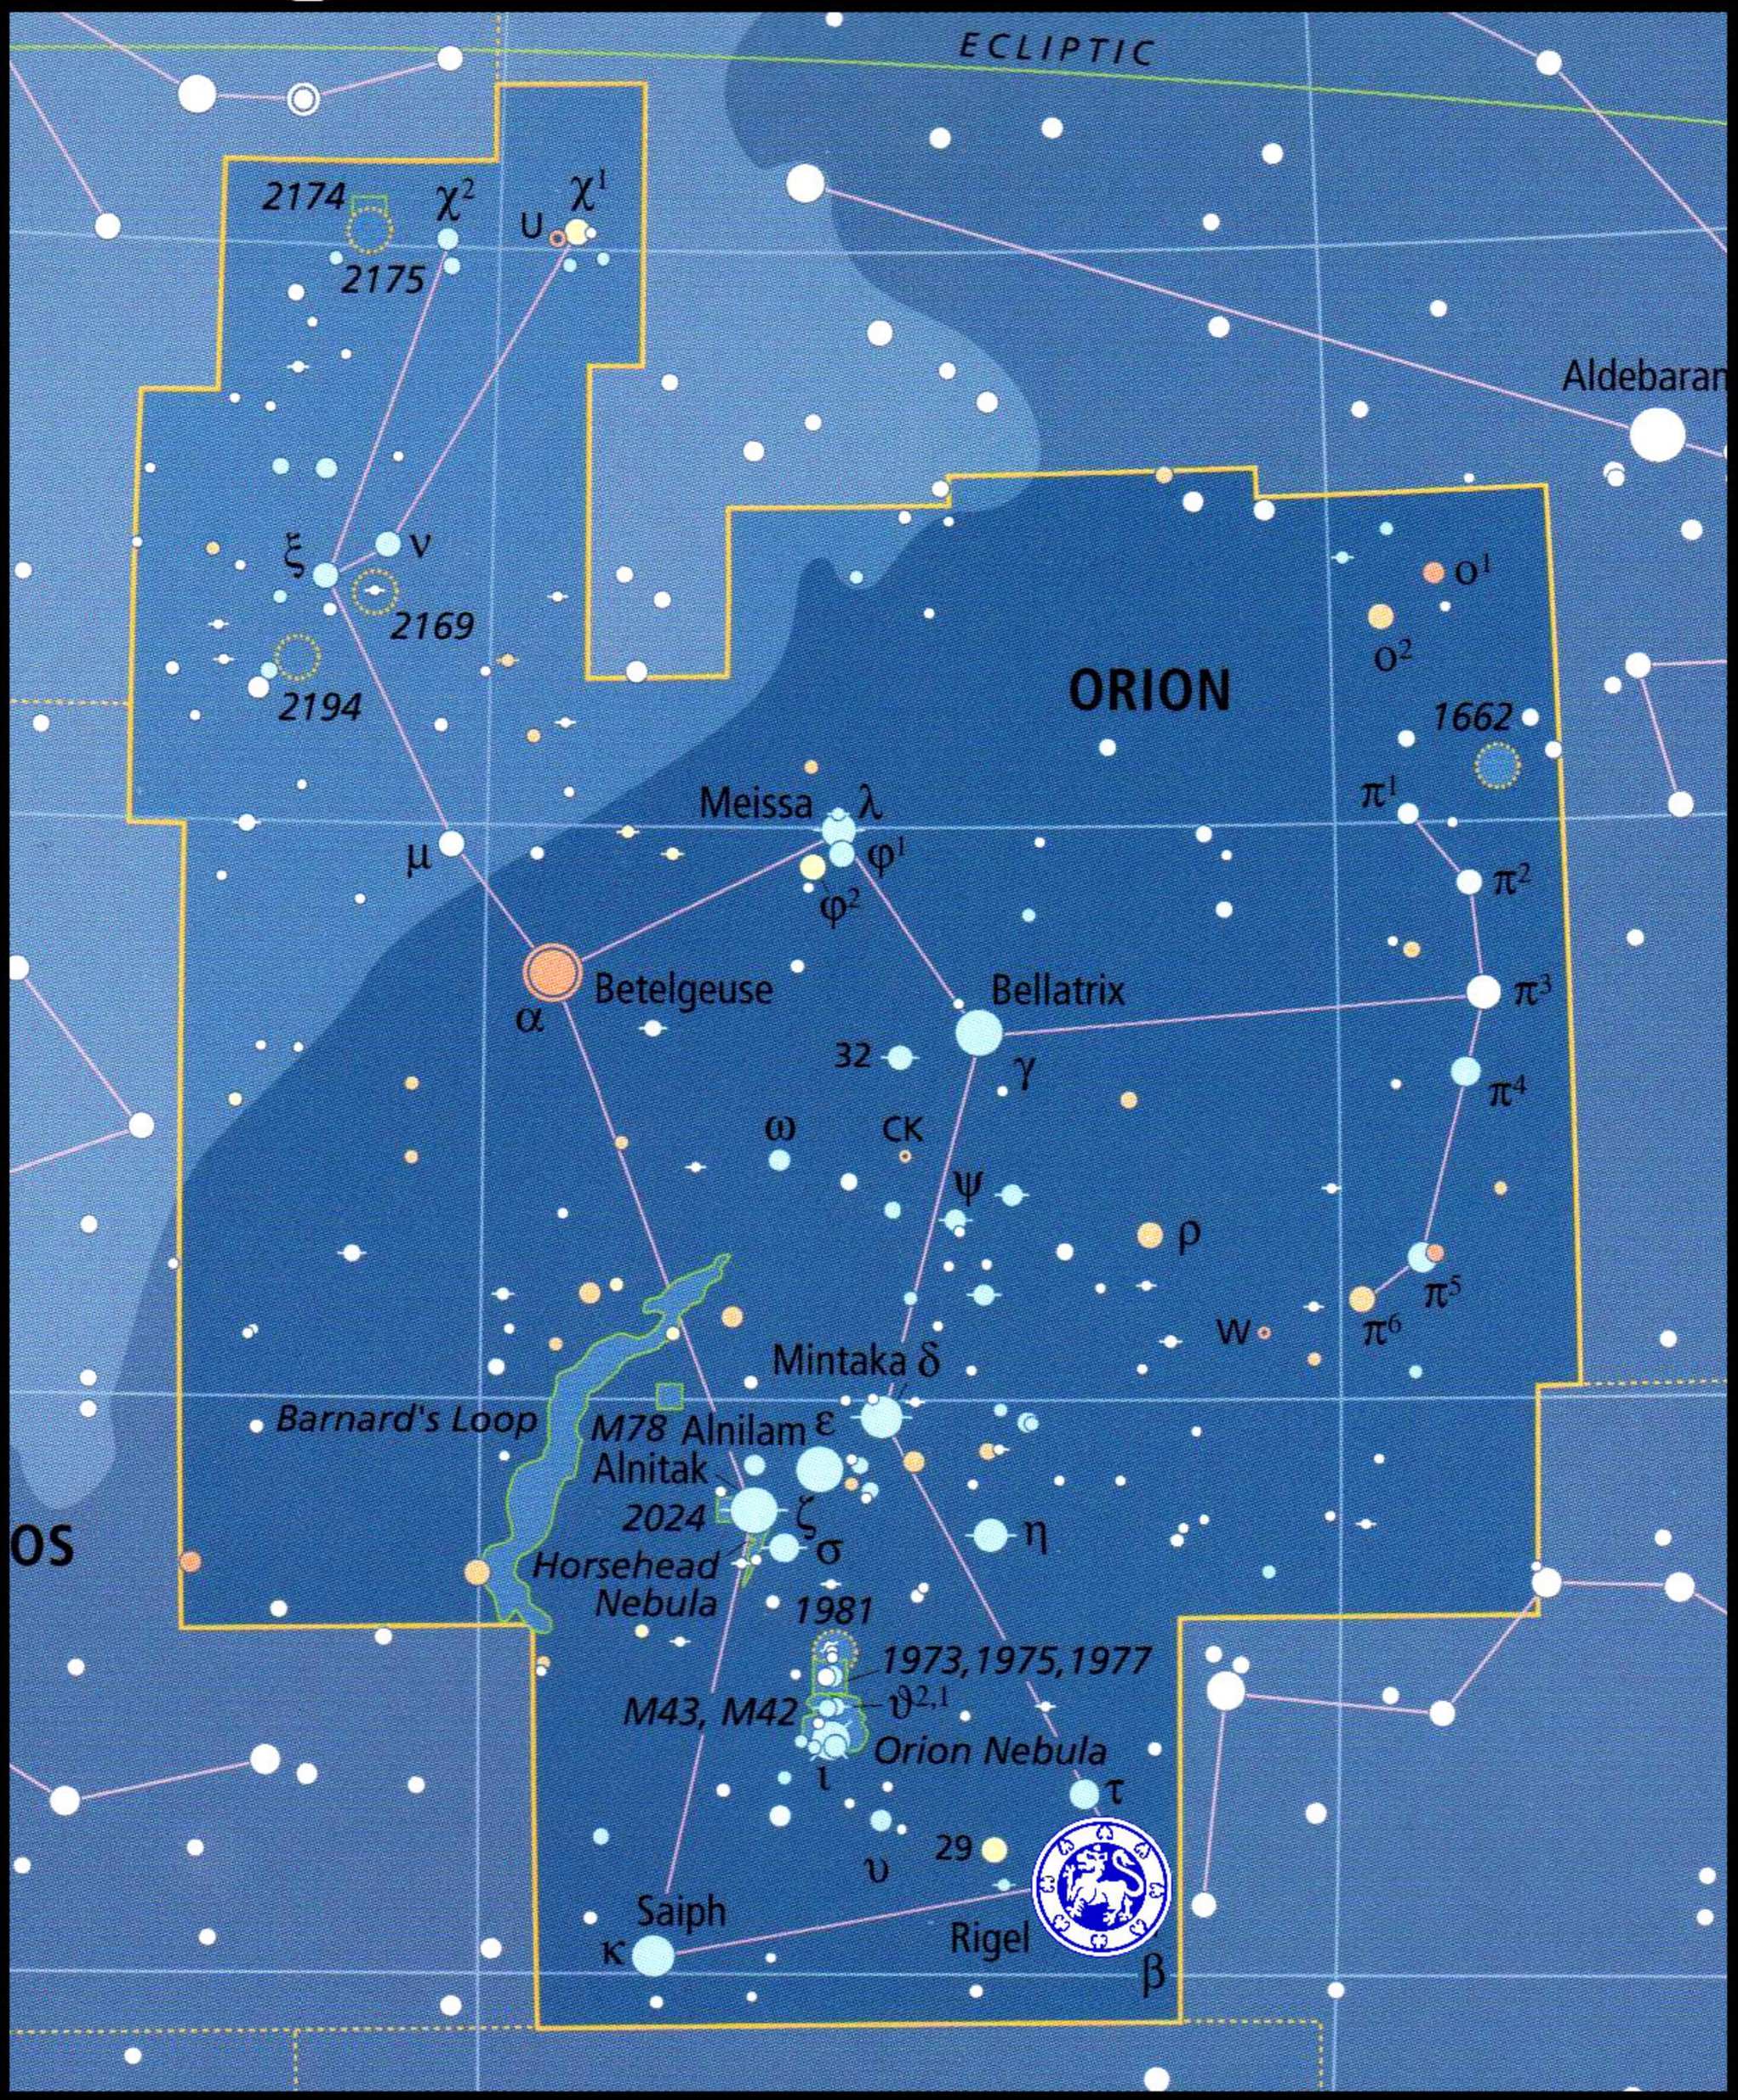

starhome